



A FASTER Pi 5!
Overclock your brand new Raspberry Pi 5!



FILE VAULT
Build the perfect backup machine



POINT 'N' CLICK
Create retro LucasArts-style adventure games

LINUX FORMAT

The **#1** open source mag



LIVE DISTROS
Instant booting for rescue, repair & more!

LEARN LINUX!

Lose the bloat, refresh your PC and embrace a clean system that delivers a truck-load of software!



PLUS: HOW TO

- » Browse and work in privacy with Tails
- » Improve your SSH security with keys
- » Get the latest Raspberry Pi OS Bookworm

CODING BBC BASIC

Relive your retro coding days on the classic BBC

CONVERT DOCS

Scan and OCR your printed documents

DIY TERMINAL

We start to code the LXF Shell from basics

LXF January 2024



Advertisement

Car, home and travel insurance plus broadband,
breakdown cover and more. Get more info or
compare quotes, for over 40 products, at go.compare



**WE'VE GOT
THE OPTIONS.
YOU MAKE THE
CHOICE.**

**GO.
COMPARE**



LINUX FORMAT



» MEET THE TEAM

We're trying to entice people to learn to use Linux, so what one piece of helpful advice would you give Linux newbies? (Apart from RTFM!)



Matthew Holder

My advice is to pick a well-supported distro, such as Ubuntu, which has a friendly and knowledgeable community to help in case of trouble. Then build your knowledge before attempting other distros, such as

Arch. And there's nothing wrong with dual-booting at the start.



Les Pounder

My advice is to try a few Linux distros. There is so much choice and it can be difficult to learn everything. So don't. Try a few distros, get a feel for them and use that as a base from which to explore the

OS, window managers, package managers and so on.



Michael Reed

If you're in the early stages of getting used to Linux, consider installing it in a virtual machine on top of your current OS. Good, free virtualisers include *VirtualBox* and *VMware*. Once you're comfortable with a

distro and know your way around, install it for real.



Nate Drake

If you're coming from Windows, choose a distro with a more traditional desktop interface, like Linux Mint. Mint also offers multimedia extras during install for smooth playback. If you have a fave Windows app,

research alternatives, such as *GIMP* to replace *Photoshop*.



Nick Peers

The best advice I can give is not to fear the terminal – you don't need a computing degree to use it. Often, firing off a terminal command is much faster than pointing and clicking your way through various dialog

boxes, and it's amazing how quickly you can pick things up.

Fresh start



As we dive into 2024 and a whole new year, isn't it time you tried something new and exciting? There seems to be a steady tide of readers who have moved from Windows to Linux, but not so much because Linux distros have suddenly become so much better or easier to use – they've been incredible productivity engines for decades now and there's a good argument that the latest

iteration of Windows apes the Linux desktop more than ever.

It seems the reason why most people move is the bloat surrounding Windows, which it simply can't escape due to its corporate roots. Adverts, tracking, pre-installed nonsense, pushing Microsoft services and products, enforced hardware upgrades... Not that Linux distros manage to avoid all of these annoyances, it's just never near to the same extent. Meanwhile, gaming – one area Linux struggled with for years – is starting to work seamlessly and sometimes better on Linux thanks to Valve's Proton, Vulkan and open source drivers.

That's besides all the class-leading productivity, server, dev and maker software that's Linux-first. We're covering just a tiny sample this issue, such as classic adventures, *WordPress*, OCRing documents, live distros, BBC Basic, C++ shell coding, Tails privacy and loads more. Linux: try it, you just might enjoy it!

Neil

Neil Mohr Editor
neil.mohr@futurenet.com



**Subscribe
& save!**

On digital and print:
see page 16

Contents



REVIEWS



Intel Core i9 14900K 19


Fast, powerful and a bit boring is not how to describe **Jacob Ridley**, but Intel's new chip fits the description nicely.

OSGeoLive 16.0 20

Uncharted no more – **Nate Drake** maps out the array of geospatial tools on offer in the specialist distro OSGeoLive.

GhostBSD 23.10.1 21

Roll over Slimer! **Nate Drake** explores GhostBSD and finds that, like Casper, it's extremely friendly and easy to manage.



Kubuntu 23.10 22

Nate Drake explores the lavish new Plasma desktop in the latest Kubuntu. Is this the greatest KDE-based distro to date?



Counter-Strike 2 24

It's CS:GO Jim, but not as we know it, says **Rich Stanton**, as he takes the esports stalwart's successor for a spin.

LEARN LINUX!

Finally had enough of your Windows PC? **Nick Peers** reveals how to road-test – then switch to – a user-friendly Linux distro. See page 32!



CREDIT: Magictorch

ROUNDUP



Live distributions 26

Michael Reed checks out five distributions that you could carry around with you on your keychain to give you Linux goodness at a moment's notice.

IN DEPTH



Tall Tails 50

Nate Drake explores how to configure the latest version of the Debian-based Tails distribution for maximum security and convenience.

CREDIT: Getty Images/Stockphoto

PI USER

- Pi news** 41
Introduced by **Les Pounder**, who is thrilled with the Raspberry Pi 5 but still keeping hold of his Pi 4. Plus, Arm buys a slice of Pi.
- Pi OS Bookworm** 42
Les Pounder has a Raspberry Pi 5 and isn't afraid to use it, now that he has the latest Raspberry Pi OS.
- Creality K1 Max** 43
Always a big fan of fancy, **Denise Bertacchi** has taken a shine to the K1's bigger, shinier brother.
- Take your Pi 5 to the next level** 44
Les Pounder has got his hands on a Raspberry Pi 5 and wants to show you how to get the best from it.
- Bake your own PC backup** 46
Christian Cawley needs to back up his PC. So, of course he chooses to do it with a Raspberry Pi.



CODING ACADEMY

- Write a Linux shell from scratch** 90
Never one to shy away from the difficult, **Ferenc Deak** takes us by the hand and helps us code a shell – from scratch.
- Build a smart-home data app** 94
Matt Holder investigates how to take data from an API and display it in a GUI for fun and possibly profit!



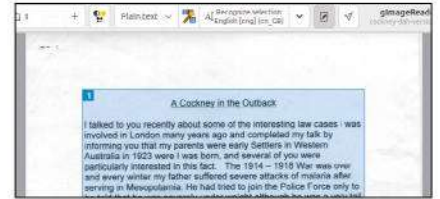
REGULARS AT A GLANCE

- News** 6
Valve console Steams ahead; foundations laid for new High Performance Software Foundation; Canonical cloud creation; industry insider opinion; plus a look at the latest distros.
- Kernel watch** 10
- Answers** 11
Neil Bothwick tackles our readers' computing conundrums concerning blank browsers, keyboard confusion, disobliging dependencies, group gripes, directory difficulties, and more.
- Mailserver** 14
Readers point out our photo faux pas and share their superior knowledge.

- Subscriptions** 16
Get your monthly Linux dose, grab some EarFun Air S1 earbuds and save money!
- Back issues** 64
Get hold of previous *Linux Format* editions.
- Overseas subscriptions** 65
Get *Linux Format* shipped around the globe.
- HotPicks** 83
Mayank Sharma has spent so much time dreaming about a Pi 5, that he won't waste more waiting for his to arrive, and will find new apps instead: *HiFile*, *Textadept*, *Kooha*, *Warehouse*, *Brave*, *Feh*, *Dosage*, *Bugdom 2*, *OpenArena*, *Linux Assistant* and *GridPlayer*.
- Next month** 98

TUTORIALS

- TERMINAL: Safeguard your secrets** 54
Shashank Sharma gets a taste for 007's life by reading fiction and sending encrypted messages to his brother.
- TESSERACT: Perfect OCR** 56
Nick Peers reveals how to extract editable text from images and printed materials with the help of OCR software.
- WORDPRESS: Final touches** 60
Michael Reed concludes his overview of setting up a WordPress site with some finishing touches and a look at plugins.



- BBC BASIC: Retro coding** 66
Turn your home system into a BBC Micro by following **David Bolton**'s advice to download BBC Basic for classic coding.
- AGS: Point-and-click adventure** 70
Nate Drake invites you to relive the glory days of point-and-click adventure games by creating your very own.

ADMINISTERIA

- Administeria** 74
Stuart Burns laments the lack of Linux on ARM, plus shares tips for improving SSH login security and for using *Docker*.
- Oxylabs** 76
Having tested more than his fair share of proxy services, **Mayank Sharma** finds this one is a breath of fresh air.
- Private Internet Access** 77
A high-value VPN with a pile of interesting features that's keeping the nefarious **Mike Williams** out of trouble.
- No Ethernet, I expect you to die!** 78
More stubborn than a spy stereotype and just as reliable, **Darien Graham-Smith** uncrosses the network twisted pair.

Newsdesk

THIS ISSUE: Valve console Steams ahead » Foundations laid for HPSF » Canonical cloud creation

CONSOLE

Valve reveals Steam Deck OLED

While it may not be the Steam Deck 2 gamers prayed for, the OLED models offer an improved screen, battery life and networking.

On 16th November, Valve made its limited edition Steam Deck OLED model available for purchase in the US and Canada. In order to minimise the kind of disruption surrounding bots and scalpers buying up the original model, besides being based in North America gamers also need to have a Steam account “in good standing” to buy the new 512GB or 1TB models. Purchases are currently limited to one per customer per week.

Those who upgrade benefit from faster NVMe SSD storage, as well as a 7.4-inch display with HDR support. The screen boasts 1,000 nits of brightness and up to a 90Hz refresh rate (previously 60Hz). Valve claims that this new screen is the main reason why the latest model is around 5% (30g) lighter than the Steam Deck LCD. The overall weight is around 640g.

Naturally, the team at Valve has also taken this opportunity to upgrade the Steam Deck's internal specs.

Chief among these is an improved 50Wh battery. The battery itself has had a major chemical makeover for faster charging, and the LED has been changed to WRGB. The power supply cable has also now been extended from 1.5 to 2.5 metres. Resume time has been improved by 30%, too.

The OLED model has an AMD Zen 2 APU with four cores, eight threads, and up to 3.5GHz clock speed, an AMD 8 RDNA 2 CUs 1.6GHz (1.6TFlops FP32), and 16GB LPDDR5 RAM (6,400MT/s quad 32-bit channels). The new model includes a dedicated Bluetooth antenna

and a new wireless chip that supports tri-band Wi-Fi 6E.

Crucially for Linux lovers, the Steam Deck OLED is still powered by the Arch Linux-based Steam OS 3, which uses the KDE Plasma desktop environment. FOSS lovers everywhere can also rejoice in Valve's announcement that there's also now preliminary support for open source BIOS and EC firmware.

Valve has stressed that it's offering the Steam Deck OLED as an experiment to see if there's a large demand. If so, it plans to roll out this model in other regions. It has also promised



The Arch-based Steam OS 3.5.5 includes support for the Steam Deck OLED, as well as overhauled colour rendering.

OPEN SOURCE FUN AND GAMES

FOSS lovers can rejoice in Valve's announcement that there's support for open source BIOS and EC firmware.

a reservation queue if it runs out of stock of a specific model.

Those who do choose to buy the 1TB model will also get a limited edition carrying case with a printed liner.

Prices start at £479 for the 512GB OLED model. For those on a budget, the entry level Steam Deck LCD with 64GB SSD is still available for £309 while stocks last. Read more at <https://store.steampowered.com/steamdeck>.

CREDIT: Valve

SOFTWARE

High Performance Software Foundation announced

The HPSF is to inspire HPC innovation and “make life easier for high performance software developers”.

The Linux Foundation has announced the formation of the High Performance Software Foundation (HPSF). The stated aims of the HPSF are to revolutionise the high-performance computing (HPC) landscape by fostering global collaboration and driving open source solutions for an increasingly diversified set of architectures.

Its specific goals are to build, promote and advance a portable software stack for HPC, as well as address its growing ubiquity in scientific computing, digital engineering and AI.

It claims to have received investment from organisations such as the United States DoE’s Exascale Computing Project and the EuroHPC Joint Undertaking, which the HPSF will use to exploit the potential of emerging technologies, including GPUs and other compute accelerators.

HPSF will serve as an umbrella project under the Linux Foundation, providing a neutral space for pivotal projects. The initial open source projects launched by HPSF include *Spack* (HPC package manager), *Kokkos* (performance-portable programming model), *AMReX*



The HPSF website boasts a star-studded list of members including AWS, HP Enterprise, Nvidia, Intel and Los Alamos National Laboratory.

(framework for solving partial differential equations), *WarpX* an exascale computing project to manage plasma accelerators, *Trilinos* an object-orientated framework for large-scale multi-physics problems and *Charliecloud* a bring-your-own software stack for high performance computing centres. They will be managed by a technical advisory committee and follow a governance model inspired by the Cloud Native Computing Foundation. HPSF representatives held a presentation at the ACM/IEEE Supercomputing Conference (SC23) on 13th November, showcasing the foundation’s vision and goals. Learn more at <http://hpsf.io>.

CLOUD COMPUTING

Canonical launches MicroCloud

The Ubuntu publisher releases new cloud software.

Ubuntu developer Canonical has unveiled *MicroCloud*, enterprise-focused software designed to simplify private cloud deployment.

Built on Canonical’s LXD container infrastructure tool, *Ceph*, and *OVN*, *MicroCloud* emphasises ease of deployment, with Canonical’s website claiming you can “deploy your cloud with a single command”. It is aimed at minimising user effort, making it accessible for enterprises of all sizes.

MicroCloud boasts automatic security updates via *Snap*, scalability from three to 50 Ubuntu-powered machines, and is licensed under the AGPL-3.0 open source licence.

Canonical promotes *MicroCloud* as “low-touch, efficient and reliable”, positioning it as

an attractive option for organisations seeking streamlined private cloud solutions.

This latest initiative is likely a strategic move by Canonical to drive *Ubuntu Pro* and support subscriptions. There are several support tiers available, as well as per-node pricing.

Feedback has been generally positive, but the limit of 50 machines may make *MicroCloud* unsuitable for organisations with extensive computing needs. Once installed, you can start using *MicroCloud* in the same way as a regular LXD cluster. Still, the touted easy setup and maintenance presumably only apply to those organisations who stay in Ubuntu’s ecosystem.

Nonetheless, *MicroCloud* could be a powerful solution for businesses. See https://canonical.com/blog/canonical_releases_microcloud.

OPINION

WHAT’S IN A NAME?



Italo Vignoli is one of the founders of *LibreOffice* and the Document Foundation.

Many users ask why *LibreOffice* is only available as a desktop application. They are probably not familiar with the *LibreOffice* Technology concept and are looking for an app with the same name for cloud and mobile.

In the case of *LibreOffice*, the cloud and mobile apps have a different name because they are built from the *LibreOffice* source code by a company in the ecosystem and not by The Document Foundation. So, one must look for *Collabora Online* and *Collabora Office for Android and iOS*.

Not only that, but there are companies such as *Allotropia* and *Collabora* itself that release a version of *LibreOffice* optimised for the corporate world, with Long Term Support (LTS) and Service Level Agreement.

Other companies, such as *Nextcloud*, leverage the *LibreOffice* Technology development platform to offer individual productivity suites to their users. This is the beauty of open source: sharing resources – here, the *LibreOffice* source code – for the benefit of everyone.

The most important thing is not the name of the product, but the fact that the software is available, and has an extremely solid, robust and flexible base in *LibreOffice* Technology.

OPINION

CLOSING THE DOOR



Joe Brockmeier
is head of community, Percona.

“Another source-available licence is making the rounds, called the Functional Source License, or FSL for short. The tag line is “freedom without free-riding”. The goal? To delay competition with the vendor for (at least) two years. After that, the code reverts to Apache 2.0 or MIT.

This seems more palatable than some other fauxpen source licences. It has some big question marks, legally, for companies that want to adopt the software. But it's far more balanced between community and commercial.

But open source isn't about providing single-vendor advantage. Yet the discussion is tilted firmly towards letting vendors do the development and placing their goals above the larger community's needs.

Open source seems to have abandoned the idea that we come together to drive things like the Linux kernel or PostgreSQL. Instead, we're negotiating with vendors about how much control to cede (and when) instead of banding together to develop software for the common good.

While I believe the FSL creators mean well, the trend is troubling. We'd be better off with a healthy commons of open source that can be supported by a range of companies. If you can't compete, it's not open, and it's not healthy for the larger community. ”

SOFTWARE

Blender 4.0 released

Latest stable version has overhauled UI and improved support for OneAPI.



New features include animation and rigging enhancements. Nodes also have collapsible sections, making them much easier to manage.

A new major release of the free and open source 3D graphics software tool *Blender* is now available for download.

Version 4.0 has returned to a flat background and incorporates a new easy-to-read font, Inter. The 3D view now has an asset shelf, which enables users to access their asset library without opening a dedicated editor like before.

The software also boasts larger colour pickers, plus you can change the colour of wireframes in all shading modes. The overlays drop-down is now organised into General and

Mode Specific. Mercifully for digital artists everywhere, the outliner select hierarchy now works with multiple selected objects.

Since *Blender 3.3 LTS*, the Cycles rendering engine has included OneAPI as a rendering device API, with support for the Intel Arc A-series. The latest version includes better support for OneAPI on Linux machines, although full support for AMD and Nvidia GPUs is still some way off.

You can learn more at: www.blender.org/download/releases/4-0/.

CREDIT: Blender Release Notes

GNOME

Foundation gets €1 million from Germany

Gnome recognised as public interest infrastructure.

Gnome is Ubuntu's default desktop, but the Gnome Foundation also contributes towards non-Gnome-specific programs like *Systemd* and *Libsecret*.

Perhaps it should come as little surprise, then, that Germany's Sovereign Tech Fund has decided to give it a cool million. The fund “supports the development, improvement and maintenance of open digital infrastructure”.

Gnome plans to modernise the platform and support features in the public interest. <https://foundation.gnome.org/2023/11/09/gnome-recognized-as-public-interest-infrastructure/>

As well as the Gnome desktop, the Foundation contributes to other open source projects.



KERNEL


Itanium heading for scrapheap

The end of an era: Will the Linux kernel drop support for ia64 architecture?

The Linux kernel has bid farewell to ia64 (Itanium) architecture during the 6.7 merge window. Initially anticipated as the future of computing, Itanium faced performance issues, while AMD's x86-64 gained market dominance.

Despite waning interest, ia64 kernel code lingered until recently, when maintainers pushed for its removal, citing impediments to development of the EFI subsystem.

Linus Torvalds has entertained the idea of a potential resurrection after a year, provided the “infrastructure is still alive”, although it seems unlikely given recent efforts to remove ia64 support from the GNU C Library, too.



Support for Itanium was introduced in 2000 but its performance fell behind AMD's x86_64.

CREDIT: Wikimedia Commons/Notes, Gnome Foundation

Distro watch

What's behind the free software sofa?

ORACLE 9.3

Oracle Linux is an enterprise-class distro backed by Oracle and built from source packages for Red Hat Enterprise Linux (RHEL). Despite Red Hat's new source code policy, this version is still application binary compatible with the corresponding RHEL 9 Update 3 release. It ships with the newest Unbreakable Enterprise Kernel Release 7 Update 2 (UEK R7U2) and also now supports automatic reboot of systems after patches are installed through DNF. Read more: www.oracle.com/linux.



It's built from RHEL source packages and is binary compatible.

BACKBOX 8.1

BackBox Linux is an Ubuntu-based Linux distribution for penetration testing and security assessments. Like version 8, the latest release (code name Sara) is based on Ubuntu 22.04 LTS, but version 5.15 of the kernel has been updated. Other improvements include an overhaul of the desktop environment. BackBox includes enhanced hacking tools such as *Airgeddon*, a multi-use bash script for Linux systems to audit wireless networks. The bootable ISO hybrid also now has UEFI support. You can learn more at www.backbox.org.



BackBox is based on Ubuntu and designed for pen testing.

CLONEZILLA LIVE 3.1.1-27

Clonezilla is an extremely efficient, open source partition and disk cloning program similar to *Norton Ghost*. The Live edition can be booted for easy cloning of system partitions. The latest version is based on the Debian Unstable Sid repository and the Linux kernel has been updated to 6.5.8. *Partclone* has also been updated to version 0.3.27 with new read and write options. There are also numerous language file updates. Read more at <https://clonezilla.org>.



Clonezilla Live is ideal for backing up and restoring partitions.

FEDORA 39

Fedora is developed by the community-supported Fedora Project. The latest version was released exactly 20 years and one day after the very first iteration, Fedora 1. Fedora Workstation now includes Gnome 45, which means there's a new workspace switcher and better image viewer. There are also updated desktop flavours for KDE Plasma, Xfce, Cinnamon and more. Fedora 39 includes the latest version of *Inkscape*, too, which also turns 20 this year. You can find out more at <https://fedoraproject.org>.



Fedora 39 includes Onyx, a Budgie-based atomic desktop.

RHEL-9.3

As the name suggests, RHEL (Red Hat Enterprise Linux) is designed for the commercial market. The latest version includes support for both legacy and UEFI boot for AWS EC2 images. Red Hat is also now compatible with the Linux storage system Stratis. This includes advanced features like thin provisioning, snapshots and caching. The *Convert2rhel* tool has been updated to support more RHEL-derived distros to make transitioning easier. Read more at www.redhat.com.



RHEL now supports transitioning from derivative distros.

OPINION

DRIVER'S TEST



Faith Ekstrand is an engineering fellow at Collabora.

“ NVK, the open source Vulkan driver for Nvidia hardware in Mesa, is now an officially conformant implementation of the Vulkan 1.0 API on Nvidia Turing hardware. You can even find it on the Khronos website under Conformant Products. It is the first time any Nouveau driver has got the Khronos conformance badge on any API.

What does this mean? Practically, it means that we can pass the entire Vulkan conformance test suite. From the Khronos perspective, it means NVK now meets the bar required to claim to support the Vulkan API officially. (There are legal implications to this that matter to the Mesa project, but most users don't care about them.) For users, it means the driver should pretty much work on Turing (GTX 16xx, RTX 20xx) and later GPUs. There will still be bugs, but those bugs are likely to be app-specific. Most stuff should just work.

We're not far off being able to advertise Vulkan 1.3 but it's all compiler work between here and there. There is also a small group of developers working on adding Maxwell support to the new back-end, so we should see improved hardware support soon, too.

NVK is still labelled “experimental” within Mesa, but the future is looking bright!

OPINION

GETTING TOGETHER



Jon Masters is a kernel hacker who's been involved with Linux for over 22 years.

“I recently attended the 2023 Linux Plumbers Conference (LPC). This was my first in-person major conference in a few years, and it was absolutely amazing to see so very many familiar faces. Plumbers is definitely the modern-day spiritual successor to older conferences where many of us first met at the turn of the millennium. It's an event “by the community, for the community”, and it retains a grass roots feel.

My main interest was Confidential Computing. I witnessed key maintainers like Greg Kroah-Hartman in the session dedicated to the subject, which included different architecture communities. It was amusing how similar their approaches all were, except for the mainframe person, who had very different opinions about how to handle device attestation and trust. I also enjoyed conversations around emerging architectures, and sessions related to bringing newcomers into the community. As many of us have been attending these events for multiple decades already, it's an important topic, so I was impressed that all of the sessions were recorded and on livestream. I would like to think this helps make such events more broadly inclusive, and you can find out more and links to the archives at <http://lpc.events>. ”

Kernel Watch

Jon Masters keeps up with all the latest happenings in the Linux kernel, so you don't have to.

Linus Torvalds announced the release of several Linux 6.7 RC (Release Candidate) kernels, noting in his -rc1 announcement message how big this release already is relative to many “historically big releases”. The main reason for this is the removal of support for the ia64 (Itanium) architecture that we've covered previously and people knew was coming for some time. Among the other contributions to 6.7 are support for the (highly opinionated) bcachefs (“the COW [Copy On Write] filesystem for Linux that

reliable, but to be able to read it even long after interest in the original filesystem (or version thereof used for a particular image) has gone. Like other contemporary filesystems, bcachefs implements compression and full filesystem encryption, as well as having built-in checksumming. Unlike some others, it uses COW (Copy On Write), meaning that updates take the form of copies of the original, so that data is stored more in a log-structured way. The idea is that file data is never destroyed, merely versioned, with older copies eventually being recycled to free up capacity for new writes.

Linux 6.7 adds the ability to disable emulation of 32-bit system calls, but this doesn't mean that 32-bit support is completely going away on x86. That topic has come up from time to time, mostly at the distro level, with some ceasing to do releases targeting 32-bit-only

“The ability to disable emulation of 32-bit system calls doesn't mean that 32-bit support is going away.”

hardware (as opposed to running 32-bit legacy binaries on 64-bit systems, which remains a thing). Still, with Linux 6.7, it will be possible for the first time to selectively determine whether a particular program is allowed to use the 32-bit Linux system call interface. The idea is that this improves security by reducing the overall attack surface, as most software doesn't need to use these legacy interfaces all the time. **LXF**

won't eat your data”), and the ability to disable emulation of 32-bit system calls on 64-bit x86 systems. The addition of bcachefs has been many years in the making. Filesystems are not generally added to Linux willy-nilly. They must be carefully engineered, well maintained and supported for many years – users generally not only want their data to be

hardware (as opposed to running 32-bit legacy binaries on 64-bit systems, which remains a thing). Still, with Linux 6.7, it will be possible for the first time to selectively determine whether a particular program is allowed to use the 32-bit Linux system call interface. The idea is that this improves security by reducing the overall attack surface, as most software doesn't need to use these legacy interfaces all the time. **LXF**

» ONGOING DEVELOPMENT

There continue to be significant developments on the Rust front. In addition to newer language requirements and subsystem support in Linux 6.7, upstream is currently discussing a big change to the Android “Binder” IPC mechanism under the topic “Setting up Binder for the future”. In addition to rewriting the core of Binder in Rust, the proposal is to leverage this rewrite to reduce overall complexity, shorten the size of “thousand line functions”, and enhance the overall security story. The authors have “left the binderfs filesystem component in C. Rewriting it in Rust would be a large amount of work and requires a lot of bindings to the file system interfaces.”

At the Linux Plumbers Conference (LPC) mentioned previously, there was the 2023

Linux Foundation (LF) Technical Advisory Board (TAB) election. Out of about a thousand eligible voters, around 200 voted in the election. The results were that several well-known kernel developers – including stable kernel maintainer Greg Kroah-Hartman and Linux Weekly News editor Jonathan Corbet – were added to the TAB for the next two years. The TAB provides technical guidance to the Linux Foundation and handles policy violations on LKML (Kernel Mailing List).

Speaking of the mailing list, several of the kernel sub-team lists are in the process of migrating servers (due to the deprecation of Mailman service by Linux Foundation) so if you are subscribed, it's worth watching to make sure you aren't missing out on emails.

Answers

Got a burning question about open source or the kernel? Whatever your level, email it to answers@linuxformat.com



Neil Bothwick
made a new-year resolution to stop talking in Python code.

Q Blank browser

On page 85 of LXF307 is an item about the *Min* browser. I would like to use it, so I downloaded the DEB file. When I ran it, I found that it needed some extra files. These were **gconf2**, **gconf-service** and **gconf-service-backend**. I installed these files and made a fresh start. The program loaded and provided its icon.

When starting there are two options: Start Browsing or Take A Tour. I took the tour, which worked OK, but when I started the browser, all I got was a blank screen. Please help me to get it working properly! I am using Ubuntu 23 LTS.

Bryan Mitchell

A Ubuntu 23.04 and 23.10 are not LTS releases, or was this a typo and did you mean you are running Ubuntu 22 LTS? We tried with 22.04 (LTS) and 23.04 with the same results. The problem appears to be that *Min* was installed on the first attempt, but not set up correctly because of the missing **gconf** dependencies. The solution was to run:

```
$ sudo apt -fix-broken install
```

to clean things up, then:

```
$ sudo apt install gconf-service
```

This pulls in the other **gconf**-related dependencies. Then reinstall the *Min* DEB

package with:

```
$ sudo dpkg -i min-1.29.0-amd64.deb
```

Min should work correctly now; it does here. If you still have problems, try running it from a terminal with:

```
$ min
```

This will not magically make it work, but it should provide some messages that indicate what is going wrong. When we installed the **gconf** dependencies but did not reinstall *Min*, this is what happened, with complaints about the GPU. After reinstalling, the errors went away and *Min* worked as you would expect.

Q Keyboard confusion

We had three PCs: a desktop, a recent laptop and an older ThinkPad. Each is set up with Linux Mint 21 and the same keyboard layout: Preferences > Keyboard > Layouts > English(US) > Options, with a third-level key defined (Left Win) for use with the Esperanto accented letters. They were all working fine with *Firefox* and other programs. Recently, we bought an NUC mini PC and set it up the same but...

Using Linux Mint 21, *Firefox* version 118.0.2 won't accept input with the Left Win key – the cursor does not move. Every other program works OK. I added EN (US, intl, with dead keys) and it allows

me to input Esperanto and other letters in *Firefox*, but only using dead keys, not using the third-level key. In the meantime, I am using the *Brave* browser. It is excellent – thanks for the review (LXF301).

What is *Firefox* doing? Is there some setting I can change to make it accept third-level input?

Donald

A Have you tried using the keyboard from the desktop with the NUC, in case it is some strange mapping of the keyboard? We have seen this on laptops where the Fn key works in reverse to how you'd expect, requiring digit gymnastics to perform certain functions. If so, this can usually be fixed with judicious use of *Xmodmap*, although this only helps if the problem occurs with all programs.

Next, check that there isn't some strange mapping within *Firefox*. Type `about:config` into the location bar, accept the warning, then type `ui.key.menuAccessKey` in the search bar. The setting for `ui.key.menuAccessKey` should be 18 – another value could be the cause of your issues.

Finally, it has been reported that with some keyboards, pressing Windows+Shift once fixes this issue, although you have to do this each time you start *Firefox*.

Q Group gripes

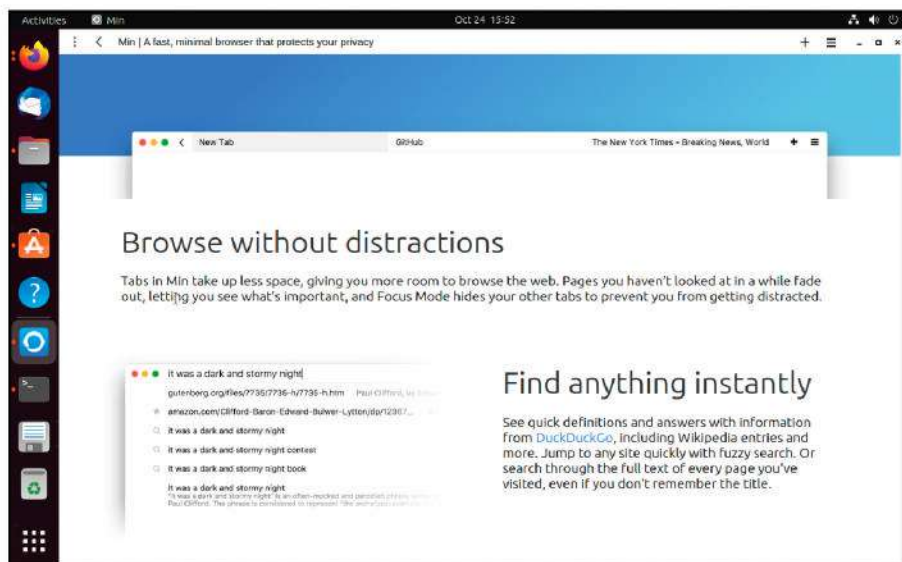
One subject that trips me up regularly is security and permissions. I often find myself in the position where UserA can't read/write files or directories created by UserB; even though UserA and UserB are in the same group, and the security bits for Group are 110 or 111 in both the directory and files. This arises often when UserB is my regular account and UserA is an app that, when I run it, runs as UserA. Such as *Jellyfin* or *Emby* servers. My usual solution is to:

```
$ chmod -R 777 /path/to/problem/directory
```

and sometimes:

```
$ chown -R UserB:UserB /path/to/problem/directory
```

as well. This seems to work but I believe such practice is frowned upon by Linux



The *Min* browser is nice, but can play up if dependencies are not installed first.

gurus who know better than me (but don't/can't explain why or what one should do to get things to work as required). One of the media drives is formatted as NTFS, a hangover from my Windows days – does this matter?

Andy Bond

A Using NTFS for anything that makes significant use of Linux permissions is asking for trouble, as you are trying to shoehorn one system's properties into a very different system. Setting permissions to 777 does work, because you are saying anyone can do anything, effectively removing any permissions settings. Note that some programs will not run if their configuration files have too loose permissions.


Linux permissions are quite a blunt instrument. They work well when trying to restrict access by user, but can't always cope when you want more than one user to be able to write and modify files in the same directory. The most effective solution is to use ACLs (Access Control Lists), but these deserve a full article of their own, so we'll look at how to manage the situation with normal permissions.

Even if you give group write access to a directory, the default behaviour is for all new files to be only writable by the user that created them. This is controlled by the umask of the filesystem, which normally defaults to 022. This is subtracted from the 777 for directories and 666 for files to give the behaviour described. A solution is to mount a filesystem with a different umask, but that statement hints at the restrictions here – the umask is applied filesystem-wide.

We recommend you create a separate filesystem for media files, then mount that with **umask=002** added to the options in **/etc/fstab**. Once created and mounted, but before you add any files, make the *Jellyfin* user the owner of the filesystem:

```
$ chown jellyfin: /path/to/mount
```

When multiple users access a program's data, group, permissions can be tricky to maintain.



The colon following the username tells *chown* to also change the group ownership to the user's main group. If you want to use a different group, add it after the colon. Now any user that is a member of that group can read, write and delete files in that directory, even those created by another member of that group.

Keeping journals

I need a logging system for error debugging and tracing with the Systemd journal. The options in **journal.conf to configure the log storage path are:**

Storage=persistent, log files will be stored in **/var/log/journal**, on disk.

Storage=volatile, log files will be stored in **/run/log/journal**, on RAM.

I would like to use the Storage=persistent option but want to store log files on disk with another path, such as **/usr/log/journal. Is this possible or do you have another suggestion?**

Sean Glover

A You need to ask what you are trying to achieve by using a different path. The standard locations in Linux are **/usr** for files that do not change – executables,

libraries, documentation and so on – and **/var** for dynamic data. Almost all daemons store their data in **/var**. If you are short of space on **/var**, be aware that *Journald* takes this into account when storing journal files. There are two options in **/etc/systemd/journald.conf**: **SystemMaxUse** and **SystemKeepFree**. They limit the use of the filesystem on which **/var/log/journal** resides. By default, the journal does not use more than 10% of the total filesystem size and leaves at least 15% free. You can change either of these to suit your needs.

It is not possible to change the location for persistent storage – unless you want to modify the source and recompile it. If you really do need to put the journal on a different filesystem, you can mount it at **/var/log/journal** or either of its parent directories. This is particularly handy if you are using volume management, such as LVM, btrfs or ZFS – then you can control how much space you leave for the journal.

Another option is to use a symbolic link from **/var/log/journal** to your chosen location. This should work but could cause a problem if the filesystem you are linked to is not mounted when *Journald* starts. This should not be the case with

» A QUICK REFERENCE TO... THE SYSTEMD JOURNAL


Most distros use *Systemd* as their service manager, but it does a lot more than manage startup services. One of the extra features is that it replaces the traditional system log with a journal. This is indexed for faster searching, which means it is no longer a pure text file that can be read with a pager like

Less, or searched with the likes of *Grep*. Instead you use *journalctl*, which has useful options. For example, **\$ journalctl -p err** only shows journal entries with a priority of error or higher, useful when trying to find problems, while, **\$ journalctl -b** shows journal entries since the last boot. You can

follow this with a number to select a different boot: **1** shows the first boot in the journal, **-1** shows the last but one, **-2** the one before, and so on. This can be combined with **-p err** to show only errors for that boot. You can also use **\$ journalctl -f** to follow the journal in real time. The **--since** and **--until**

options let you view entries over a specific time period, while **-t** shows entries with a particular identifier, such as **\$ journalctl -b -t sudo** to show all uses of *sudo* since the last boot. There are a lot more options, but you can still use the old tools, like *Grep*, if you prefer:

```
$ journalctl -b | grep somestring
```

■ Birdtray depends on Thunderbird, but it is possible to fool it into installing without it.

a separate mount as above, as *Systemd* will not start the persistent journal until the storage location is mounted. So think carefully before using a symlink, and try not to break the filesystem hierarchy standard by putting dynamic data in */usr* – standards are usually there for a good reason. For example, some distro installers wipe */usr* when installing over an older release, but leave data directories like */var* and */home* untouched.

Q Replacement bird

I have Linux Mint 21.2 installed and am now using *Betterbird* instead of *Thunderbird*, so I want to get rid of the latter. My problem is that *Apt* won't let me uninstall *Thunderbird* without uninstalling *Birdtray*, because it has *Thunderbird* as a dependency. And if I let it uninstall *Birdtray*, I won't be able to reinstall it because *Thunderbird* won't be there.

But I like *Birdtray*! I simply changed its command-line setting to point to *Betterbird*'s executable and it works great. I could use *Betterbird*'s built-in tray icon option but it is not as useful. Is there a way to change that dependency?

Madeleine Carr

A There is a way to do this, by creating a dummy package for *Thunderbird* that doesn't actually install anything. First you need to install the *equivs* package:

```
$ sudo apt install equivs
```

Then run *equivs* to create a control file for the package you want:

```
$ equivs-control thunderbird
```

This creates a text file named *thunderbird* in the current directory. Load it into a text editor and make sure it contains the information you want. Once you are happy, create a package with:

```
$ equivs-build thunderbird
```

This creates a *thunderbird-dummy* DEB file in the current directory that you can install with *Dpkg*, but first you need to uninstall the real *Thunderbird* package. As you say, this would require you to uninstall *Birdtray*, too, but you can reinstall it after installing the dummy. It is not possible to uninstall a package without removing packages that depend on it using *Apt* or any of the graphical package managers, but you can do it with *Dpkg*:

```
$ sudo dpkg --remove --force-depends thunderbird
```

Now install the dummy with:

```
$ dpkg --install ./thunderbird-dummy*.deb
```

Use the actual name of the DEB file and reinstall *Birdtray* as necessary. You may need to repeat the process of generating and installing the dummy package at some point in the future if *Birdtray* depends on a later version.

Q Home from home

I am using Ubuntu 22.04.1 on two machines. I have a backed-up *home* directory from computer A – say *homeA.bak*. Computer B already has a *home* directory for the principle user (let's call him userB), as set up upon installation.

Now I want to create a new userB without a *home* directory and then create a *home* directory for userB that will be *homeA.bak*. I assume from what I read that *home* directories are not like regular directories, in that they contain other files that make them *home* directories – you just can't make a directory called *home* and copy backed-up files to it and have it act like a *home* directory. Is this correct and how should I proceed?

Nicholas Reed

A There is nothing special about a *home* directory. It is simply a directory that is owned by a particular user

and recorded as their default home in the */etc/passwd* file. By convention, *home* directories for regular users are */home/username*, but this doesn't have to be the case. System daemons often have *home* directories in */var*, such as */var/www* for a web server, */var/spool/cups* for the CUPS print daemon, and so on.

In your situation, it is usually best to use the standard distro tools to create the new user, including their *home* directory at */home/userB* in your case. You could copy the contents of the backup there, remembering to copy all hidden files – those whose name starts with a dot. But it is probably easier to delete and rename:

```
$ sudo rmdir /home/userB
```

```
$ sudo mv homeA.bak /home/userB
```

This has to be done as root, using *sudo*, because you are dealing with directories owned by different users. Also, userB cannot be logged in or you can't delete their *home* directory. Then you need to set the ownership of the new *home* directory and all of its contents, to the new user.

```
$ sudo chown -R userB: /home/userB
```

The *-R* operates recursively, changing the directory and everything in it. The colon following the username is shorthand that tells *chown* to also change the group of everything to userB's group. Because Linux filesystems use numeric IDs to store ownership information, this is necessary even if you use the same username on both systems. Because the user was the first created on the original system, they have a user ID of 1000, but it is 1001 for userB because the existing user is already using 1000, so this step is required. **LXF**

GET HELP NOW!

We'd love to try to answer any questions you send to answers@linuxformat.com, no matter what the level. We've all been stuck before, so don't be shy. However, we're only human (although many suspect Neil is a robot), so it's important that you include as much information as you can. If something works on one distro but not another, then tell us. If you get an error message, please tell us the exact message and precisely what you did to invoke it.

If you have, or suspect, a hardware problem, let us know about the hardware. Consider installing *hardinfo* or *lshw*. These programs list the hardware on your machine, so send us their output. If you're unwilling, or unable, to install these, run the following commands in a root terminal and send us the *system.txt* file, too.

```
uname -a > system.txt
```

```
lspci >> system.txt
```

```
lspci -vv >> system.txt
```

Mailserver

WRITE TO US

Do you have a burning Linux-related issue that you want to discuss? Write to us at *Linux Format*, Future Publishing, Quay House, The Ambury, Bath, BA1 1UA or email letters@linuxformat.com.

Punched in the face!

You should know better! The image attached to the Letter of the Month in **LXF308** is not of punched cards at all. It is of folded paper tape.

I have attached an image of real punched cards, attributed to use on a DEC PDP9. I do not recall ever working on that model – yes to almost every one of the PDP8 and 11 families, at least one PDP15, every variant of VAX and most Alphas.

John

The photo is, of course, paper tape. It appears to be six-track, which I've never seen. I have used seven and eight-track, even five-track in the early days (1964).

David Harris

Neil says...

Thanks for the correction John, David and others – you're obviously absolutely correct. The photo we used doesn't even look like cards. Not sure what's going through my brain...

Terminal Format

I have a great deal of sympathy with A Confused of Ross-on-Wye. I have beaten him on age, being 89, rising 90 if I'm spared. I was given a car-boot-load of cardboard boxes back in 1990, and told, "This is your computer." I found it was an Amstrad, I think, and fortunately the connecting plugs were such that you could only really set it up one way. This I did, switched on and began to learn how to use a computer.

Now I, too, have given up on Windows and run my machines on Linux Mint Cinnamon or LMDE. I can download an ISO, make a bootable USB stick and install the OS without difficulty. I can and do install the programs I want from the software manager or *Synaptic Package Manager*. Most of the time I can get along quite nicely, but am often stymied by the words "tar", "gz" or similar. I know I should venture into the terminal – but what to do when I get there? Who will supply the necessary hieroglyphics to make it work?

Linux Format is good for those who speak the language, but for the uninitiated, something like *Computeractive* would be good. I tried to get it to stop being 'Windowsactive' and to include Linux stuff – but failed. Is there a *Terminal for Dummies* book?

Ian Sheppard

Neil says...

I guess there's an expectation (in the Linux world) that people will just 'pick this up' as they go along. Looking around, there's not many good beginner guides; there's lots of guides for developers or administrators, and, of course, the overly meticulous Man system itself.

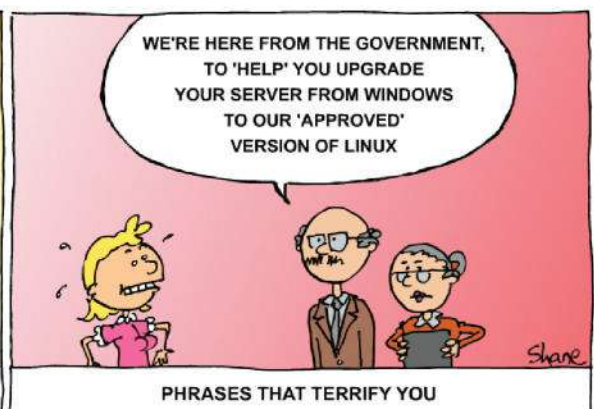

See <https://linuxjourney.com> for a modern resource or a more formal guide at <https://linuxupskillchallenge.com>. An older classic is the online game <https://web.mit.edu/mprat/Public/web/Terminus/Web/main.html>.

From a magazine perspective, we'd have to cover the same beginner topics over and over. It's almost like having to explain how to use icons and menus each issue. I'll pull together our old terminal guides into a single PDF. Anyone can download this from:



That's not a punch card. *That's* a punch card – on an actual loom in the Manchester Science and Industry Museum.

Helpdex





» LETTER OF THE MONTH

Scraping the barrel

My heart sank when I noticed LXF309 had a tutorial on web-scraping in Python. Yes, sure enough, another article on using Requests and Beautiful Soup. Both are fine libraries, but together they do not constitute a scraping framework. The article claimed 70 lines of code; of the code that the excellent Scrapy framework gave me, I had to change only 35! You were right to caution against hitting sites too hard, but I solved that problem by uncommenting a single line to enable the autothrottle extension. Yes, Scrapy can use XPath, a powerful and portable domain-specific language for querying web pages, and it can be character-building at first, but it can also use Beautiful Soup. These days, the mere parsing of HTML is often neither sufficient or necessary for web-crawling.

The real power of Scrapy is the separation of the code that makes requests from the code that parses them. It makes it much harder for novices to make a mess of things by providing a hand-rail. For more advanced users, features like headless browsers, testing and monitoring are provided by well-integrated third-party plugins. I spent four years in the music industry crawling hundreds of thousands of accounts across dozens of sites, I know whereof I speak. https://github.com/augeas/lxf_archive

Dr Greenway

Neil says...

Thank you for your thoughtful reply, you obviously know a great deal more on the subject than myself (which wouldn't be hard) or indeed David, who admits knows nothing of Scrapy. We didn't set out to write a scraping framework as such; I see our coding tutorials as basic standalone guides to encourage people to play around and experiment with coding. So, they tend to be just examples rather than best in class. There's always more than one way to skin a cat... But Scrapy does seem to be something we should look into.

Neil says...

I suspect we'll get back to looking at Vlang at some point, but right now we're kicking off a multi-part C++ project creating a basic shell from, erm, basics... While I keep getting suggestions for fun Python projects, like web scraping. It's so hard to fit it all in! LXF



It's V for Vlang, I'm not sure that's how the old saying goes?



The non-trademark infringing Tux can be seen here wielding a generic laser sword.

https://bit.ly/lxf_terminal, or Mayank Sharma's Master the Terminal feature from https://bit.ly/lxf222_feat.

Vlang

In terms of modern compiled languages, you often refer to Go and Rust, but I was wondering when you will be touching on Vlang again. I find it very easy to use and update. For example, you can update Vlang simply with `$ v up`.

The error reporting is comprehensive, including suggested corrections. The fact that it compiles to native code makes it one of the fastest to run, which appeals to the real-time area that I work in.

You did a brief introduction, but version 4 is now released and has some significant changes and enhancements. Further insight and examples in *Linux Format* would be interesting and might help give some impetus to the project, and help avoid it falling by the wayside through lack of following.

Ken Shail



shane_collinge@yahoo.com

SUBSCRIBE Save money today!

SUBSCRIBE

today and get a pair of

EarFun Air S1 Earbuds

Worth £69.99!

Included with every print subscription.

EarFun Air S1 Earbuds

- » 30 hours' play time with the charging case.
- » Personalise with the *EarFun Audio* app.
- » Multipoint connection function for easy switching between various devices – connect to two devices at once, be it a mobile phone, tablet, smartwatch or PC.
- » Active noise cancellation.
- » Intuitive touch control and all-in-one volume control – a feature missing from 95% of true wireless stereo earbuds.
- » Wireless charging and USB-C charging.
- » Sweat and water resistant.
- » Secure and ergonomic fit.
- » Ability to activate Voice Assistant.
- » Among the most technologically advanced headphones we've ever offered. Available with every print subscription. Get yours today!



SUBSCRIBE NOW!

www.magazinesdirect.com/LIN/XE92

Call **0330 333 1113** and quote **XE92**



» **PLUS:** Exclusive access¹ to the *Linux Format* subs area!

Thousands of DRM-free PDF back issues and articles! Get **instant access** back to issue 66 (May 2005) with tutorials, interviews, features and reviews. At **linuxformat.com**

¹ Only available to www.magazinesdirect.com subscribers

» CHOOSE YOUR PACKAGE!

One payment. A whole year of issues delivered straight to your door or device. Simple!

ANNUAL PRINT EDITION



Only **£62.99**

13 issues of *Linux Format* in print over 12 months, **saving 25%**

ANNUAL DIGITAL EDITION



Only **£51.99**

13 issues of digital *Linux Format* over 12 months, **saving 42%**

OUTSIDE THE UK?



Turn to **page 65** for more great subscriber deals!

^{*}Terms and conditions: **Offer closes 31/01/2024.** Offer open to new UK subscribers only. Pricing is guaranteed for the first 12 months and we will notify you in advance of any price changes. Please allow up to six weeks for delivery of your first subscription issue (up to eight weeks overseas). Your gift will be delivered separately within 60 days after your first payment has cleared. Gifts only available to subscribers on the UK mainland. Gift not available with a digital subscription. Gifts only available in the colour black. The full subscription rate is for 12 months (13 issues) and includes postage and packaging. If the magazine ordered changes frequency per annum, we will honour the number of issues paid for, not the term of the subscription. For full terms and conditions, visit www.magazinesdirect.com/terms. For enquiries please call: +44 (0) 330 333 1113. Lines are open Monday-Friday 9am-5pm UK time or email help@magazinesdirect.com. Calls to 0330 numbers will be charged at no more than a national landline call and may be included in your phone provider's call bundle.

Winter Savings

SUBSCRIBE

FROM JUST £6



PLUS OFFERS ON BINDERS, GUIDES & SPECIALS



See the entire range online at

MAGAZINESDIRECT.COM/WINTER23

or phone **0330 333 1113**

Outside of the UK? See our international offers online

Offer closes 31st December 2023. Offer open to new subscribers only. Direct Debit offer is available to UK subscribers only. We will notify you in advance of any price changes. All gift subscriptions will start with the first available issue on-sale after December 2023. If you would like your gift subscription to start with an earlier issue, you can choose an earlier issue when ordering or you can contact customer services. Orders purchased for yourself will start with the next available issue - please allow up to 6 weeks for delivery (up to 8 weeks outside of the UK). Payment is non-refundable after the 14 day cancellation period unless exceptional circumstances apply. For full terms and conditions, visit www.magazinesdirect.com/terms. For enquiries please call: +44 (0)330 333 1113. Lines are open Monday-Friday, 8:30am-7pm, and Saturday, 10am-3pm GMT (excluding Bank Holidays) or e-mail: help@magazinesdirect.com. Calls from the UK to 0330 numbers will be charged at no more than a national landline call, and may be included in your phone provider's call bundle.

Intel Core i9 14900K

Fast, powerful and a bit boring is not how to describe **Jacob Ridley...**

SPECS

Socket: V LGA1700
Process: Intel 7 10nm
Cores (P+E): 8+16
Threads: 32
P-cache: 640KB L1, 16MB L2, 36MB L3 (shared)
P-core: 3.2GHz (5.6GHz boost, 6GHz TVB)
E-cache: 1.5MB L1, 16MB L2, 36MB L3 (shared)
E-core: 2.4GHz (4.4GHz boost)
Unlocked: Yes
GPU: Intel UHD 770
GPU clock: 300MHz (1.65GHz max)
Ex units: 32
Display: 4, eDP 1.4b, DP 1.4a, HDMI 2.1
Mem max: 192GB, ECC support, 2-channel
Mem speed: DDR5 5,600MT/s, DDR4 3,200MT/s
PCIe: v5 or v4 20-lanes
Processor base power: 125W
Maximum turbo power: 253W

For the fastest chip ever to grace a gaming PC, the Intel Core i9 14900K is about as boring a CPU refresh as one could be. We've seen it all before. The massive core count, the architecture, the 6GHz clock speed; the 14900K might do some things better than the 13th gen but it rarely matters for much. That doesn't make it any less of an awesome chip – it is – it just doesn't push the needle forward much, if at all.

That's why PC builders have a tough decision ahead of them: do you buy a 13th-gen chip and save a couple of quid or invest in a shiny new 14th gen for the often intangible performance benefit and bragging rights?

The 14900K comes with the same recommended customer pricing as its predecessor, the Core i9 13900K, at launch. That's £579, give or take. The big question for any PC builder is whether the 14900K's distinct brand of bland is worth spending extra money on top versus the Core i9 13900K?

What's new with the Core i9 14900K? Not all that much. As the flagship processor of a refreshed Raptor Lake generation, it is almost a carbon copy of the Core i9 13900KS. You're looking at a 24-core processor, split between eight Performance-cores (P-cores) and 16 Efficient-cores (E-cores).

Graphics grumble

The iGPU present on the 14900K is the UHD Graphics 770. Unlike the shiny new Arc Xe-LPG iGPU found on Intel's new Meteor Lake mobile processors, this iGPU doesn't pack the punch required of it for 1080p gaming in lieu of a discrete graphics card.

You can use both 600-series and 700-series motherboards with this processor. Naturally, this means you can stick with a DDR4-compatible motherboard with the 14th gen. This will surely be the last chip generation to make this offer, however.

It's actually in power draw where things become a little more interesting for the 14900K. It has a lower processor base power at 125W compared to the 13900KS's 150W, despite being marginally faster. The P-cores and E-cores are both faster at max turbo: an increase of 200MHz on the P-cores and 100MHz on the E-cores.

On that front, the 14900K delivers smooth gameplay throughout our testing. You could make a similar case for AMD's 3D V-Cache gaming CPUs, however, because there are games that favour either Intel or AMD's top chips in this metric. There are a few titles where AMD still has Intel beat in terms of gaming performance, too, though there's usually little between these rival processors.

The 14900K shows its best performance in non-gaming benchmarks. It's mostly in the more traditional benchmarks that you see the benefit of



Spot the difference between this and the Core i9 13900KS.

mildly increased clock speeds. It's the fastest desktop CPU we've ever run through the X264 encoding benchmark, and it nearly knocks AMD's Ryzen 9 7950X3D off the top spot in *Blender*.

The 14900K is pretty dominant in *Cinebench R23*, and when you start throwing unlimited power at this chip, it'll give you a little more in multithreaded tests. It's not all that convincing, however, as a good run on a juiced-up 13900K had it beat once.

It's the performance in these sorts of intense workloads that you need to care about to justify buying the 14900K. Otherwise, just grab a 14600K – it's seriously almost as good for gaming and still a great runabout for multitasking and editing.

We've been spoiled with a couple of great chip generations back to back, and this sort of refresh just doesn't really inspire us to wax lyrical about how wonderfully adept it is at this, that or the other. Besides a couple of neat tricks, it feels as though we've seen this chip before – it's the Core i9 13900K dressed up in a monocle and top hat. **LXF**

VERDICT

DEVELOPER: Intel
WEB: www.intel.com/processors
PRICE: £579

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	9/10	VALUE	7/10

With no tangible benefit over the often cheaper Core i9 13900K, stick with whichever Core i9 you can find cheapest.

» **Rating 8/10**

OSGeoLive 16.0

Uncharted no more – **Nate Drake** maps out the array of geospatial tools on offer in the specialist distro OSGeoLive.

IN BRIEF

OSGeoLive is an excellent distro for specialists who regularly work with geospatial data. As it's based on Ubuntu, it could be suitable for everyday use with some tweaking.

SPECS

CPU: 1GHz
Mem: 1GB
 (2GB recommended)
HDD: 20GB
Builds: i386, AMD64

OSGeoLive is an OS that allows users to deploy a wide variety of open source geospatial software. In other words, this is a niche distro that will chiefly interest scientists and explorers.

It's sponsored by the OSGeo Foundation, a non-profit that supports geospatial open source software development.

The OS itself is based on Ubuntu, though the number of bundled apps means that the live ISO is over 4GB. There's also an optional preconfigured VMDK disk image for virtual machines.

As the name suggests, the distro works perfectly well in live mode but more recent versions of OSGeoLive can also be installed to a hard disk with 20GB or more of space.

You may want to do this for the sake of security, as the default root password is 'user'. Similarly, passwords for the bundled apps are listed in a plain text file on the desktop, though in fairness they are only for programs that work with publicly available open source data.

If you decide to proceed with setup, you'll notice that the installer is no different from Ubuntu's – indeed, the splash screen actually uses the name of Ubuntu rather than OSGeoLive. One quirk we did notice was that the installer defaults to German, though you can click Deutsch in the drop-down menu easily enough to switch languages.

As it's based on Ubuntu, this is a fully-fledged operating system with access to Ubuntu repositories, so in theory it could be used as your main OS if you have use for the bundled geospatial software.

This would require some technical chicanery, though, because Ubuntu's *Muon* package manager is noticeably absent, as are the usual tools you'd need for an OS suitable for day-to-day use, such as an office suite and media player. Still, this isn't the intended purpose of OSGeoLive, which focuses on allowing you to store, view, edit and analyse spatial data of all kinds.

As hefty as the ISO is, it's clear there's been some effort to cultivate only the best Linux tools. Upon login, you'll see the desktop contains a series of folders that are used to break these tools down into categories such as Navigation And Maps and Spatial Tools.

For those who are new to geospatial software, there's a shortcut to a folder full of sample data, such as maps and detailed images. There's a number of free maps, and tools such as *JOSM* (Java OpenStreetMap Editor) have helpful wizards and welcome pages for those using them for the first time. Some of the tools



OSGeoLive categorises the various geospatial mapping and data analysis tools into folders on the desktop, which contain app shortcuts.

are standalone but others are designed to be run in the default *Firefox* browser (currently version 116.0.2).

One notable web app is *Jupyter Notebook*, which can create and share documents that contain live code, equations and visualisations. It can be used for data cleaning and transformation, numerical simulation, statistical modelling and even machine learning.

Monteverdi is one of the more useful standalone apps. It's an open source visualisation tool that works well at data compression. It takes advantage of the image processing tools in *OTB* (*Orfeo Toolbox*).

The latest version of OSGeoLive has also undergone some pruning. For instance, we were unable to find *Mapnik* when searching via the app launcher. This open source toolkit for developing mapping applications was present in previous iterations of OSGeoLive.

The same is true for *MapSlicer*, an app for creating graphical overlays of standard maps such as those in OpenStreetMap. Since both tools are still in active development, we assume the OSGeo community had its own reasons for removing them from the distro. **UX**

VERDICT

DEVELOPER: OSGeo Foundation

WEB: <https://live.osgeo.org>

LICENCE: LGPL v2.1

FEATURES	7/10
PERFORMANCE	8/10

EASE OF USE	8/10
DOCUMENTATION	9/10

OSGeoLive is a niche distro that does one thing very well in offering a Swiss army knife of precision geospatial tools.

» **Rating 8/10**

GhostBSD 23.10.1

Roll over Slimer. **Nate Drake** explores the latest GhostBSD and finds that, like Casper, it's extremely friendly and easy to manage.

IN BRIEF

GhostBSD began as an OS to hack Gnome using FreeBSD tech. Although it's now switched to MATE, it still lives up to its reputation of offering a user-friendly way to explore BSD.

SPECS

CPU: 1GHz
64-bit
Mem: 4GB
(8GB recommended)
HDD: 15GB
(30GB recommended)
Builds: AMD64

According to the official website, GhostBSD is “a simple, elegant and friendly BSD operating system for desktops and laptops based on FreeBSD”. As the website diplomatically puts it, “BSD is generally considered beyond the average computer user’s knowledge.” Those niche users who want to try BSD but struggle with text-only installers seem to be GhostBSD’s intended audience, as it combines FreeBSD with its own graphical tools for ease of use.

GhostBSD has been in active development since 2010 and previously used Gnome. (The name is a portmanteau of ‘Gnome hosted on FreeBSD’.) The official desktop environment is now Mate 1.26.0, although there’s also a community edition using Xfce.

You can obtain the most recent version (23.10.1) from the main website both as a direct download and via BitTorrent. Should you fire up the 2.7GB ISO in live mode, you’ll see there’s a simple graphical installer.

The installer leverages OpenZFS, making it easy to install GhostBSD on ZFS alongside other operating systems. On the flip side, unlike FreeBSD, there’s no automatic support for full-disk encryption.

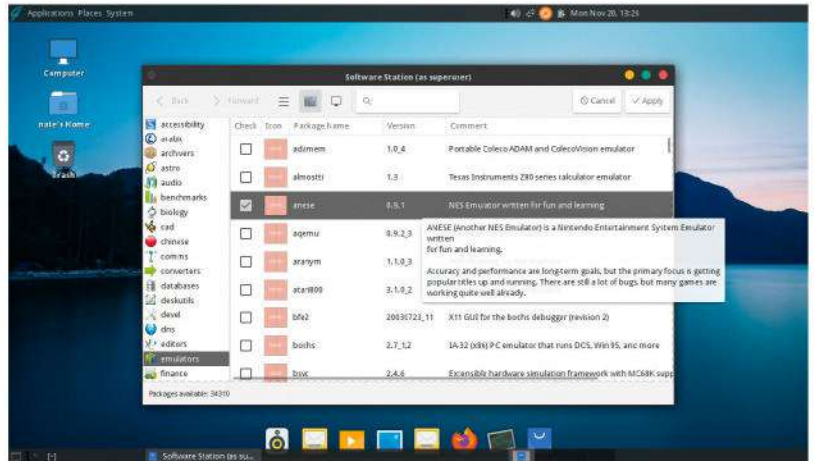
This is a shame, because the installer itself notes that one of the intended goals of GhostBSD is to offer an OS that’s secure and protects your privacy. Still, compared to other BSD distros we’ve reviewed, setup was an absolute breeze and install was complete within four minutes. Although we had issues setting a screen resolution higher than 1024x768 in live mode, once install was complete, GhostBSD’s VESA driver immediately offered all available resolutions.

In a nod to its FreeBSD roots, at boot users are still presented with an ncurses text menu offering options like being able to boot to single-user mode.

We used the official build with the MATE desktop, so after logging in we decided to start with tweaking its look and feel. GhostBSD includes 12 wallpapers. The bundled theme is customised but you can switch to Vimix, which is available in both light and dark versions.

The OS follows a slow-rolling release model for the sake of stability. This means you’re unlikely to see bleeding-edge versions of apps. Still, the bundled programs are few but very well thought out.

Although it didn’t load by default, we went to Applications > Accessories > Plank to get a handy dock. Here we discovered GhostBSD ships with version



Software Station enables you to explore and download apps via category or search bar. Hover over an application for a description.

118.0.2 of the *Firefox* browser, as well as the latest release of *VLC Media Player* (3.0.18 Vetinari).

Media playback is handled by *Rhythmbox* (3.4.6), which is designed for Gnome but can open MP3 files without issue. Email is handled by *Evolution* (3.4.4), developed by the FreeBSD Gnome Team. By default, pictures are handled by the *Eye of MATE* image viewer.

If you want to explore other available packages, you can do so via the integrated *Software Station*. This graphical app is similar to *Synaptic*, enabling you to explore categories on the left or search for specific apps via the bar at the top. The look and feel aren’t as polished as, say, Gnome software, but it’s functional.

Once installed, you can configure default and startup applications via the menu in Preferences.

Overall, it’s clear the developers have done everything they can to make setup and use of a BSD operating system as user-friendly as possible. We can’t find any fault with GhostBSD itself, though newcomers to open source might struggle with being unable to run proprietary software like *Google Chrome*. **LG**

VERDICT

DEVELOPER: GhostBSD Team

WEB: www.ghostbsd.org

LICENCE: FreeBSD Licence/GPL

FEATURES	8/10	EASE OF USE	9/10
PERFORMANCE	9/10	DOCUMENTATION	8/10

GhostBSD lives up to its claim to make BSD OSes user-friendly. Setup takes minutes and can be easily configured.

» **Rating 8/10**

Kubuntu 23.10

Nate Drake explores the lavish new Plasma desktop in the latest Kubuntu. Is this the greatest KDE-based distro to date?

IN BRIEF

Kubuntu continues to be an excellent choice for KDE lovers. The installer is easy to set up. Although not all Gear apps are bundled, they are simple to download via Discover.

SPECS

CPU: 2GHz dual-core
Mem: 4GB
HDD: 25GB
Builds: AMD64

Kubuntu has been in development since 2004 as an official flavour of the Ubuntu operating system. Current iterations, including the most recent version 23.10 (Mantic Minotaur), use the KDE Plasma desktop.

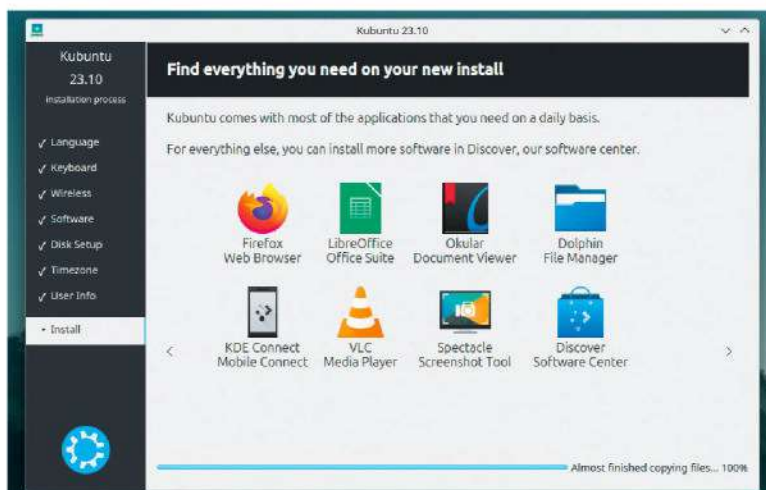
Given that it has an Ubuntu base, Kubuntu 23.10 will be supported for nine months. Crucially, 2023 is the last year that Kubuntu will support 32-bit users, as it is currently only available for AMD64 architectures.

If you choose to download the hefty 3.7GB ISO, you can take it for a test spin in live mode or install to a drive. We chose to do the latter and, as we'd expect from a Ubuntu-based OS, the install wizard was a breeze, offering the option to install third-party codecs, partition the drive and, most crucially, encrypt the system drive via LUKS. The experimental ZFS install options are not available by default, though power users may be able to configure these manually.

As it was loading, the installer also helpfully displayed colourful icons of Kubuntu's default apps and prompted us to use software centre Discover to download more as necessary. It also provides helpful links to Kubuntu's extensive official support pages and community forums.

The official release notes state that Kubuntu 23.10 ships the eighth scheduled bug-fix release of KDE Plasma 5.27 (5.27.8).

KDE Gear has been upgraded to version 23.08, which means updates to a large number of tools and software libraries. Amongst these updates is one that hides temporary files generated by applications in the *Dolphin* file manager – a feature request has been pending since



The helpful installer displays the various applications you can download via Discover. Make sure you check out the default programs first.

the turn of the century! You can also open duplicates of a *Dolphin* tab just by double-clicking on it.

After reviewing the KDE Gear 23.08 release notes, we saw that some of the updated apps such as *Merkuro* (which replaces calendar) and *Itinerary* are not installed by default, but can be downloaded via Discover. It does ship with v23.08.1 of advanced text editor *Kate*, though.

After checking the KDE website, we discovered that post-install there's also a new welcome utility to assist newcomers to the Plasma 5.27 desktop. This didn't display automatically, nor could we install it via Discover, as the software centre claimed it was already installed.

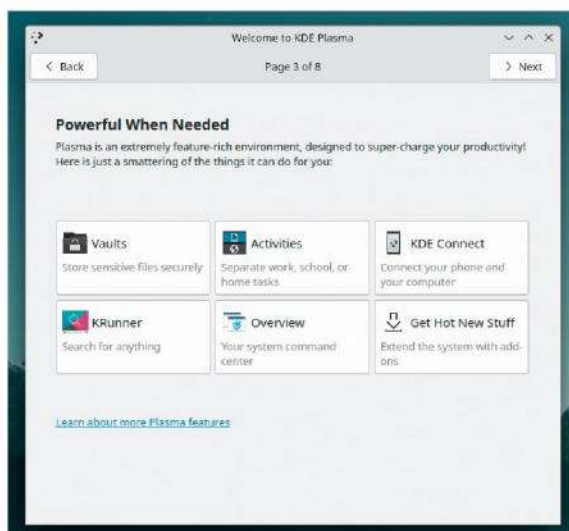
Luckily, Kubuntu includes the *Konsole* terminal utility, which we used to run `sudo apt install plasma-welcome`. This provided an excellent grounding in how to configure Plasma, pointing users towards system settings, as well as KDE-specific features such as *KDE Connect* to hook up your mobile phone. We were pleased to see that by default the option to share anonymous usage information with KDE is disabled, though the project claims to have a strict privacy policy.

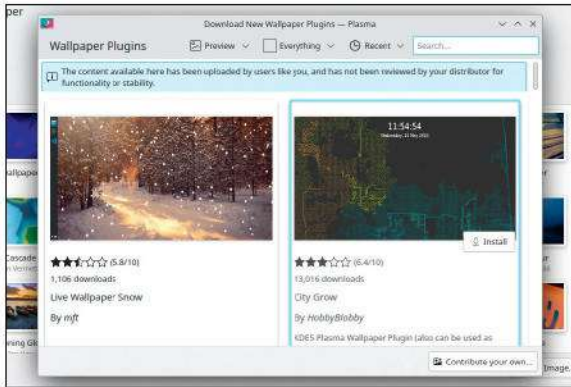
As we had *Konsole* fired up, we ran the new `kinfo` command. This displays system information such as the versions of Plasma, Frameworks and the kernel you're running, so has potential for troubleshooting.

The desktop itself deserves special mention, because if you right-click to Configure Desktop And Wallpaper, you're offered a choice of no fewer than 36 desktop backgrounds. Click Get New Plugins to install community-created backgrounds, including live images.

You can also change your Wallpaper Type via the drop-down menu. Options include Plain Color or Slideshow. We also noted two options named Hunyango and Haenau. After searching online, we discovered these categories encompass animated QML wallpapers. As a number of users have mentioned in bug reports, some

The new Plasma welcome wizard doesn't display automatically post-install in Kubuntu but you can manually install it via Konsole.





Kubuntu 23.10 comes with a number of wallpapers. Click Get New Plugins to download more from the KDE community.

description here would be useful, to save KDE newcomers from resorting to search engines.

Despite the failings of Discover to install the *Welcome* app, the layout is extremely intuitive. Categories are listed along the left, along with colourful icons for popular apps. Some of the most popular ones, such as *VLC* and *GIMP*, are displayed along the top of the window.

Another nice touch of Plasma is that if you click the application launcher and search for a program that isn't installed, the system displays a link to get it.

As for default apps, the bundled browser is *Firefox 118.0.1*. This isn't quite the most recent version (199.0.1 at the time of writing) but it does mean it includes some of the most recent enhancements, including automated local translation of web content. There's also much better browser fingerprinting protection through use of the *FDLIBM* maths library and restricting font visibility.

Kubuntu 23.10 also ships with *LibreOffice 7.6.2.1*. This includes some excellent and long-awaited upgrades, such as a page-number wizard, improved accessibility features and an overhauled help interface.

As of Ubuntu 20.04, *KMail* has given way to *Thunderbird*, which has now been updated to version 115.4.1 SuperNova. This contains numerous fixes as well as a new feature to manage newsgroup subscriptions.

The OS also comes with *Qt 5.15.10*, which includes over a hundred bug fixes and security updates.

Graphics are managed by *Gwenview*, KDE's default image viewer. If you choose to install *GIMP*, *Gwenview* can also now open the image editor's default XCF format. KDE's bundled document viewer *Okular* has also been tweaked to make digital signature options easier to use.

The *Harunda* media player can play sound files, as well as online videos via youtube-dl. Kubuntu also comes

» A TREATISE ON TILING

Kubuntu includes version 5.27 of the Plasma desktop, so now supports creating tiling layouts, as well as the ability to place windows inside the same. While this isn't as effective as a dedicated tiling window manager, it can be a good way to manage workspaces.

To get started, hit your Super key and T. This displays the default layout of two vertical tiled windows on the left of the screen with a double-sized tile on the right. From here, you can choose to split tiles horizontally or vertically, or add a new floating tile. You can also drag the edges of tiles to resize. Once you're happy with the layout, in theory you can just hold the Shift key and drag a window into its chosen tile.

By default, any windows dragged to the right or left edge of the desktop are tiled, too. You can also drag the edges of tiled windows to resize them – the other tiles change their size accordingly.

If you jump into Settings > Workspace Behavior > Screen Edges, you'll also see an option to Trigger Quarter Tiling. By default this is set to Outer 25% of the screen. In practice, we found this hard to do with the mouse but you can also use the keyboard. To activate quarter tiling, hold the Super key and press an arrow key followed by another arrow key in quick succession. To deactivate tiling, just hold Super, then tap right repeatedly until it's back to normal.

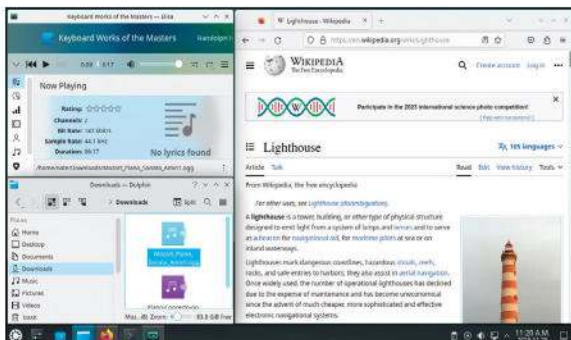
with the *Elisa* music player, though during our tests, we noted that MP3 playback defaults to *Harunda*. On the plus side, while previous reviews noted *Elisa* couldn't handle Ogg Vorbis (OGG) audio files, the version bundled with Kubuntu 23.10 played them without issue.

Elisa Shuffle and Repeat settings are also now controllable via MPRIS. This means you can control these features from the Media Player widget, or via *KDE Connect* on your phone.

Beneath the hood, Kubuntu now includes version 6.5 of the Linux kernel, incorporating a number of improvements, including performance boosts to both the ex4 and the Btrfs file systems. There's also improved forwards compatibility with the new Wi-Fi 7 and USB4 v2 standards.

When reviewing Kubuntu, it's very easy to miss all the hard work that goes into overhauling the user interface and improving apps, because the UI always feels so familiar and easy to use.

It's important, however, not to take it for granted. At first glance, Kubuntu 23.10 may not seem too different visually from the previous release, but there has been a huge number of improvements to ensure this is the OS of choice for millions of users. **LXF**



The KDE Plasma desktop now supports tiling. Hit Super+T to change the default layout and create new tiles.

VERDICT

DEVELOPER: Canonical/Blue Systems

WEB: <https://kubuntu.org>

LICENCE: GPL

FEATURES 9/10

PERFORMANCE 9/10

EASE OF USE 9/10

DOCUMENTATION 9/10

Kubuntu continues to be visually stunning and easy to set up. The new Plasma desktop is a pleasure to use.

» **Rating 9/10**

Counter-Strike 2

It's CS:GO Jim, but not as we know it, says **Rich Stanton**, as he takes the long-standing esports stalwart's successor for a spin.

SPECS

CPU: 64-bit, four-threads
Mem: 8GB
Storage: 85GB
GPU: AMD GCN+, Nvidia Kepler+, Vulkan driver
OS: Ubuntu 20.04 64-bit

Valve announced *Counter-Strike 2* at the start of 2023 and has been running it in a limited beta ever since, swapping one or two maps at a time, adding isolated features, and generally withholding the full package that it unveiled at the end of September. That limited rollout made it hard to get a handle on just how different it feels from CS:GO, but now we have no option. When CS2 went live, CS:GO became a beta branch on Steam. Despite some mixed messages from Valve, CS:GO has not been disappeared, Stalin-style, and remains accessible, albeit only with community servers.

All change, no change

Perhaps *Counter-Strike 2*'s biggest problem is that number. In a series that has always been iterative, new *Counter-Strikes* have been pitched as variants on the existing game: vanilla *Counter-Strike* is better known as *Counter-Strike 1.6* (after a particular patch version), then there's CS: *Condition Zero*, CS: *Source* and CS:GO. Whatever changes these versions made, they weren't positioned as replacements in the way CS2 has been.

That leads to CS2 initially feeling underwhelming. As one wag joked, "CS2 has only been out a day but feels like it's been here 10 years." Apart from the quality jump in maps, weapons and all visual assets, many of the improvements over CS:GO are back-end, under-the-hood things that aren't obvious to players, yet the game itself comes branded as the first true sequel in 24 years.

Our first games in CS2 proper, having played it a fair amount in beta, were nothing like moving from *Halo 2* to



Halo 3, fundamentally different multiplayer experiences. Instead it almost feels more like a director's cut or a definitive version of something you've been playing for years (because you have). We've seen a bunch of community grouching about the movement feeling lighter, and even complaints about the feel of the gunplay, but debating these very fine details is a long-standing tradition for competitive FPS players. That's not to say anyone's wrong, but having played the game for two decades, we don't detect huge differences.

Holy smokes!

The overall coat of gloss CS2 receives is, of course, very welcome. But the familiarity dulls some of that upgrade, and unless you put the two games side by side, the faithfulness means a lot of the changes aren't immediately obvious. On Valve's CS2 landing page, there are some shots of the same areas in both games that let you slide back and forth, and give some idea of what we're talking about.

The biggest mechanical difference, and you see it in nearly every match, is the smoke physics. Utility in CS is as important as being able to shoot straight, and at all skill levels is the difference between winning and losing. Smokes give control over areas of the maps for a time and can be used to block entrance routes for the attackers or to block the defending team's vision (and much more besides, of course).

While the bullet-hole effect is beautiful to look at and can provide a split-second of vision, the combination of smokes and high-explosive grenades is a genuinely new element to the game (an HE grenade disperses a good chunk of the smoke cloud for what feels like five

New weapons are promised but they've sped up the knife throw already.





I We doubt this is the room with a view EM Forster had in mind.

or six seconds before it returns to normal, granting vision through the smoke). This gives HE grenades a new form of utility, and vice versa – one decoy smoke can now bait several frag grenades from the defenders.

At its best, CS is a mind game. You can beat better players and teams by fooling them, using your grenade resources to gain ground, and by being smart about how you counter an incoming play. CS2 has added a lot of small things to the mix, some of which may in the long run prove more consequential, but this is the element that will most transform the meta. It's also the one thing about CS2 we've not seen anyone complain about.

Aces all round!

The CS2 UI is a vast improvement over CS:GO's, which while nice enough, had kinda been assembled over years. In CS2, things like the buy menu, thanks to *Valorant*'s example, are vastly improved. Small changes, such as the ability to buy both M4 models in game, are huge, and our favourite addition of all is the playing cards.

When you make a kill in CS2, you get a small playing card at the bottom-middle of your UI. Each kill in a round adds up to a max of five playing cards and, as achieving this has always been an 'ace' in CS, it fits perfectly. Does it have any great functionality? Not really. But it looks and feels tremendous, and somehow adds to the excitement of a good round that could become spectacular.

So beautiful...

CS2 is a much brighter and more colourful game than CS:GO. We love the new look of the maps and the interplay between them and the Source 2 lighting and particle systems. They are leaps and bounds over what we had before, while retaining in many cases the exact elements they had beforehand, but everything from water to fire behaves differently, and the new impact effects on gunfire, in particular, are incredibly satisfying.

A big disappointment is the absence of Arms Race, a community mode that became popular enough to be incorporated into CS:GO as an official mode. CS:GO had a whole bunch of secondary modes like this that haven't made the initial cut, but Arms Race (aka Gun Game) was our warmup: you chase kills over multiple rounds, and each kill upgrades your weapon for the next. Things like this and Team Deathmatch (Deathmatch is present) we're told will be added back in over time.

Valve has clearly focused on getting the competitive side of CS2 right as the biggest launch priority, and

things such as the new matchmaking and ranking system are polished overhauls. But competitive is also about what's around it, and what else you can do in the game when you're not queuing for ranked, and CS2 feels somewhat anaemic in this regard.

Ranked and unranked

CS2 introduces a much more in-depth stats system for Premier mode that you can view through the main menu, and it feels like a positive change but also way too early to say such a thing definitively. The various competitive games we've played so far have included some wildly

different skill levels, with some players fragging 30-plus and others only putting up single digits, but you have to put that down to this being the dawn of the first day, and give the game some time to begin putting us all in baskets. This is an entirely new ranking system from CS:GO and it doesn't look as though anything from the older one carries over, so we'll reserve judgement until it's had a few weeks to sift through the game's million-plus daily players (there are 1,441,329 players in game right now).

When you first log into CS2, CS:GO players receive two gifts. There's a CS:GO music kit, which if activated replaces CS2's audio with the CS:GO equivalents, and a CS:GO commemorative coin (these are issued for various occasions and events in game). They're not a big deal, and certainly don't have much impact on the game itself, but were a fun surprise and somewhat softened the realisation that, yeah, CS:GO was yesterday, and CS2 is today and tomorrow. But we'll always have these little reminders, and the memories. **LXF**



I A threatening hazy orange cloud emanates from this man's pointy thing.

I Ha, ha – the guy on the left forgot to bring his gun. Sad.



VERDICT

DEVELOPER: Valve

WEB: www.counter-strike.net

PRICE: Free

GAMEPLAY	9/10	LONGEVITY	9/10
GRAPHICS	7/10	VALUE	10/10

Just as the original *CS:GO* infuriated *CS1* players, CS2 will do the same, but that takes nothing from its fun and ambition.

» **Rating 9/10**

Roundup

Linux Lite » EasyOS » AntiX
» Porteus » Ubuntu



Michael Reed
has been using
Linux for so long
that he remembers
when a live
distribution booted
from a floppy.

Live distributions

Michael Reed checks out five distros that you could carry around on your keychain to give you Linux goodness at a moment's notice.

HOW WE TESTED...

As usual, we tested the distributions on a virtual machine for convenience. But we also wanted to get some insight into how well these distributions worked on genuinely lower-end hardware, so we tested the various flash drives on an old Asus laptop that we thought was typical of the type of older machine to be found in the back of an office cupboard. The specs included a 2.1GHz Intel dual-core CPU and 4GB of RAM; so, we're not talking super-low specs.

We kept an eye out for problems, but by and large, none of the distributions were slow on this hardware.

We tested all of the common functions in each and every distribution, and these included installing applications, browsing the web, and viewing and playing media. We put each distro into a mode whereby storage would be persistent and changes would be permanent, as we felt that this is the most useful way of using a live distribution.



Live distributions are distros that run directly from a medium such as flash drive without having to be installed on a hard drive. They can give you a fairly complete desktop and loads of full, useful applications. All you should have to do is put the flash drive into an available slot, turn the computer on, and press a key to enable USB booting. Live distros can get you out of a scrape when your system is messed up, and are also an excellent way of showing the uninitiated what Linux is capable of.

The Linux Lite distribution is specifically aimed at providing a desktop, along with

documentation for use, which is suitable for Windows switchers. Ubuntu should be familiar to most of you, and we've decided to make the installer ISO into a live boot system to see how the mainstream distro stood up to the more specialised distributions in this role. EasyOS is a Swiss army knife of useful tools and unconventional approaches at every turn. Meanwhile, Porteus provides a fairly lightweight desktop and the means to start adding applications that you choose. And last but not least, AntiX is a good all-rounder with some useful applications and a clean, attractive desktop.

CREDIT: Getty Images/Stockphoto

Default application selection

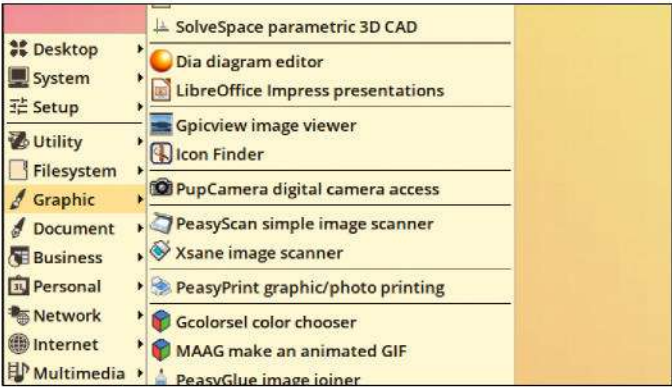
If all you have is the flash drive, what does that allow you to do?

The Linux Lite app selection is a little heavier than a standard distro loadout, with a few extra tools such as *GIMP* and *VLC*, which are full-fat apps for technically minded users. If this distribution was handed to an average office worker or a home user, they could probably scrape by through a day's work with it without having to add anything else.

Porteus gives us some basic utilities, such as a file manager and text editor, a media player and a music manager. However, there is no graphical web browser. We understand that this is a fairly bare distro, designed to be expanded and customised, but we would rather have had a browser from the start, rather than a media player. An argument could be made that this could drag in some unwanted dependency cruft, but Porteus has quite a clean expansion architecture for adding and removing software.

EasyOS includes a large selection of apps. Expect to see well-known tools such as the *Chromium* web browser, *LibreOffice* and *GIMP*, along with a large array of smaller tools that cover areas as diverse as scientific software, a virtual machine manager, graphics tools and video editing. It can feel like an eccentric muddle of lesser-known apps, but EasyOS could be ideal if you don't have much time or a reliable online connection.

The AntiX application loadout is fairly extensive. It covers typical business and home uses, such as office tasks, web



EasyOS has a mixture of mainstream and slightly quirky apps when first installed. We won't knock this selection, though, as it covers most areas of computer usage.

browsing and media playback. There is certainly more to the app offering than you would expect from a normal Linux distro, as mainstream apps such as *Firefox* sit alongside *Dillo*, a lightweight browser. The selection would get you through a typical workday.

The Ubuntu software layout covers the basics of typical office tasks, with a web browser and email application, and even a few luxuries, such as a media player. We suspect that most users would have resorted to adding some extras to the software list before a full day of work is done on this setup.

VERDICT			
LINUX LITE	7/10	PORTEUS	5/10
EASYOS 9	9/10	UBUNTU	6/10
ANTIX	8/10		

The EasyOS application selection can handle most computing contingencies, even if some of the choices are less familiar tools.

Flavours and variants

Multiple versions for different niches and requirements.

Linux Lite is the easiest of the bunch to summarise because it only has one edition, which uses the Xfce desktop environment and with Ubuntu as its base distro.

AntiX offers variants that make use of either *Systemd* or an alternative, *Runit*. The 'full' variant comes with a choice of four switchable, lightweight window managers and a range of apps. It's the recommended release, and the one we chose. There are also lighter variants that scale all the way down to one with no GUI.

Porteus is based on Slackware and comes with 32-bit and 64-bit options. There are seven variants with different desktops. We're sometimes a bit wary of distros that offer a lot of different desktops as, naturally, the application selection can never be as carefully honed. However, Porteus offers a minimalist application selection to be expanded upon anyway. There is also a kiosk edition. We chose the Cinnamon desktop version because this gives a good balance between features and looks.

Ubuntu comes in a number of different flavours, each with a different desktop and set of applications. We've decided to use standard Ubuntu and to install the ISO to a flash drive using

language/	25-Sep-2023
modules/	03-Oct-2023
updates/	20-Aug-2023
Porteus-CINNAMON-v5.01-x86_64.iso	03-Oct-2023
Porteus-GNOME-v5.01-x86_64.iso	03-Oct-2023
Porteus-KDE-v5.01-x86_64.iso	03-Oct-2023
Porteus-LXDE-v5.01-x86_64.iso	03-Oct-2023
Porteus-LXQT-v5.01-x86_64.iso	03-Oct-2023
Porteus-MATE-v5.01-x86_64.iso	03-Oct-2023
Porteus-OPENBOX-v5.01-x86_64.iso	03-Oct-2023
Porteus-XFCE-v5.01-x86_64.iso	03-Oct-2023
sha256sums.txt	03-Oct-2023

Download the Porteus variant ISO that you want from the website and go from there. The ISOs themselves are impressively small at around the 400MB mark.

MKUSB. This makes it so that changes to files in the **home** directory and any system changes are permanent.

EasyOS uses parts from Puppy Linux and Debian, but it deviates from typical Linux distros in a number of ways in terms of how it manages things like the root user, giving it less in common with the base distro than most Linux distributions.

VERDICT			
LINUX LITE	5/10	PORTEUS	8/10
EASYOS 9	6/10	UBUNTU	6/10
ANTIX	8/10		

Ubuntu has a lot of flavours, but they aren't, specifically, live distributions. AntiX, however, offers a good range for live use.

Graphical user interface

Snappy and familiar out of the box.

The whole point of these distros is to offer a comfortable, typical desktop on an ad-hoc basis. There will be a few people using a flash-drive Linux distribution as their main system, but those users are in a minority. In some contexts, you might have to hand the system over to someone who's not a Linux expert, so familiarity in terms of layout is often important.

For typical uses of a live distribution, a light, responsive desktop is more appropriate than a heavyweight one that guzzles resources. There really is no technical reason, these days, why a light desktop running on hardware dating from the last 15 years can't be both fully-featured and responsive.

Being able to carefully hone every aspect of the desktop theme isn't vitally important on a live distro desktop, but we do expect to be able to change a few things such as the colour scheme and the font size, to aid clarity if nothing else.


Linux Lite7/10EasyOS7/10

The desktop, provided by the Xfce desktop environment, uses a grey and white colour scheme. It's fast, with a clear overall look. If we tell you that there is a bar across the bottom of the screen with a task launcher, a task switcher and a status area, you're probably picturing the standard desktop layout that most people expect to see, and you'd be exactly right.

The launcher is searchable and pops up when the super key is pressed. When searched, application descriptions are taken into account.

Like a few other Xfce-based desktops, the resize area at the sides and corners of windows is thin, making it difficult to grab.

Thunar, the file manager, is a good choice in terms of speed, familiarity and functionality. The desktop doesn't break any new ground in terms of how it looks or operates, but there are advantages to an arrangement that nearly every computer user can recognise at first glance.




Like most of EasyOS, the desktop aesthetic is unconventional, making gratuitous use of pink and yellow. That's not to put it down, but it's far removed from the greys and blacks we're used to.

The lightweight JWM window manager provides the desktop, and the layout is standard, with a taskbar and launcher menu at the bottom. We were a bit disappointed that the launcher isn't searchable or launchable with the super key, but long lists of apps do have divider bars indicating subcategories.

The interface makes quite a lot of use of icons on the backdrop, and like most of this desktop, the icons are monochrome, contrasty and extremely clear. Icons are launched with a single click, and this is also true of ROX, the file manager.

EasyOS makes some slightly unusual choices with the components that make up the desktop, but it is simple, fast and extremely clear-looking.



Software installation

The options to add applications and expand the system.

EasyOS boasts four package management systems. SFS is a compressed format that is never installed, intended for large applications and the system files themselves. PKGget is a format that is closer to a normal Linux package. EasyOS can also use the more mainstream ApImage and Flatpak formats. Each format has its own GUI.

Ubuntu on a live flash drive uses the same formats and tools as regular Ubuntu. This means that it can download packages from the Ubuntu repositories, and it has Canonical's Snap system built in for more up-to-date packages. Both of these have a GUI front-end.

Linux Lite is connected directly to the Canonical software repositories and its own repository, and can use the same installation tools as Ubuntu. Synaptic is included as a package management front-end. In addition, Linux Lite has its own software installation front-end that includes about 30 common applications, such as Audacity and Firefox. Unfortunately, during our testing period, updating the Linux Lite repository failed with

an error, even though it was possible to install software. Software from the Canonical repository software installed as expected.

AntiX has its own package manager connected to Debian and MX Linux repos. From the command line, Apt works as expected.

Porteus has its own system of modules. For example, there is a front-end to install a browser, and this installs a module file. Unfortunately, some manual file copying is required to activate this module and make it persistent. We got the feeling that the modules system is designed for customisation of a system for possible distribution, rather than casually adding software.

VERDICT

LINUX LITE	6/10	PORTEUS	7/10
EASYOS 9	7/10	UBUNTU	8/10
ANTIX	7/10		

EasyOS feels fragmented but, as ever, it covers a lot of bases. Ubuntu and AntiX are both solid solutions.


AntiX**6/10**

AntiX works differently from the others as it comes with a number of combinations of file manager and window manager, and these can be switched between freely without restarting. We'll concentrate on the default setup which makes use of *IceWM* and the *zzzFM* file manager.

This uses the conventional layout with the standard taskbar at the bottom, but we were disappointed the launcher wasn't searchable. Some useful, real-time system information is displayed on the backdrop.

ZzzFM is a slightly unconventional file manager. The problem with including it is that for live boot scenarios, familiarity is important, and there are lots of lightweight file managers that provide that.

The default window manager and theme combo has a white-on-grey colour scheme, with good contrast and clarity, giving a clean, fast desktop that has some unusual choices, which might hold a live distro back if used in an emergency.

**Porteus****7/10**

We chose the edition of Porteus that comes with Cinnamon, a desktop that uses modern Gnome components but presents a more traditional layout. It's not quite a lightweight desktop, but it's snappy on most hardware and it offers a full range of modern features.

The taskbar is at the bottom of the screen in the usual place. The launcher is double width, with a column for categories and another for the actual launch icons. This makes sense because there's not much point in saving screen space when the user activates the launcher. It is searchable and launched when the super key is pressed.


It's not bristling with innovative features, but we feel that *Nemo* is a good choice as the file manager for a live distro. It does everything that a typical user is likely to need and it's obvious how it works at a glance. We could say the same for the desktop as a whole.

**Ubuntu****7/10**

Ubuntu has flavours that cover most of the major desktop environments. In our case, we chose standard Ubuntu. This means that the desktop is Gnome 3. The layout here is a vertical bar along the side of the screen. It's a sensible arrangement because it makes more sense to lose screen width rather than height.

The colour scheme is the standard Ubuntu one, which makes use of white and grey with a maroon backdrop. There are certainly no problems with clarity.

The application-launching mechanism won't be as immediately familiar to as many people as the classic one with the taskbar at the bottom, but it's easy to learn and it's efficient once you know it. Generally, you press the super key or click on Activities in the top-right corner. You can then start typing the name of an application. There is an alternative mode, which is more like the icon palette of the Android operating system.



Performance and compatibility

Plug in and play, or plug in and start praying?

All of the distributions that we tested seemed well suited to running on slower hardware.

AntiX didn't want to boot on our old laptop until we added the **acpi=off** kernel parameter in the GRUB bootloader. Following that, AntiX booted to the desktop, but couldn't set the screen resolution correctly, because the correct modes weren't available. In the virtual machine, hardware detection worked correctly. Once booted, running **free -h** in a terminal showed 390MB in use, which is pretty impressive.

EasyOS wasn't perfect in the area of hardware detection either. We were never able to get sound output to work on the old laptop. It looked as though the sound was playing, but no sound actually came out of the speaker. On another laptop, the trackpad was detected, but we had to press down unusually hard. Memory usage was 224MB on boot. EasyOS felt even snappier than the others in most areas.

Porteus detected all hardware correctly and used 481MB when booted from the live flash drive.

Linux Lite also found all of the hardware correctly on the laptop. It's not a huge deal, but it used a bit more RAM than some other distros at 649MB on a fresh boot.

Ubuntu, equipped with Gnome 3, running in live mode used 1.2GB of RAM, which is quite a lot more than the others, and it might be worth bearing in mind if you intend to use your flash drive on low-end, older hardware. It wasn't slow, but we felt it wasn't quite as snappily responsive as the other desktops when running on the laptop. All hardware was detected and configured.

VERDICT

LINUX LITE	8/10	PORTEUS	8/10
EASYOS 9	7/10	UBUNTU	7/10
ANTIX	6/10		

Linux Lite and Porteus sit in the middle in terms of frugal memory usage and correct hardware detection.

Persistent storage

These distros aren't as useful if they can't remember changes you make.

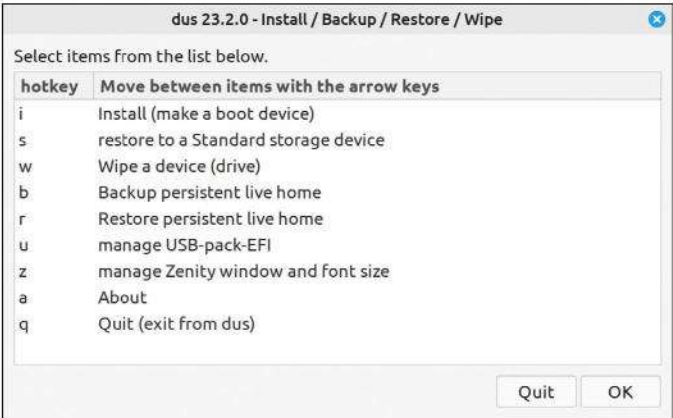
Making permanent changes is rather simple on EasyOS. The first time it boots on a flash drive, it expands to fill the drive for storage for files and changes to the system. From then on, you click on the Save icon on the backdrop if you want to make any changes permanent. You are also prompted to save every time you shut the system down.

For Ubuntu, we followed the official Ubuntu recommendation and used **MKUSB** to create the live flash drive, and this gave us an operating environment that was much the same as the standard Ubuntu.

The official instructions for Linux Lite also suggested that we use **MKUSB**, and that's what we did. This splits the partition between data storage and storage for system changes. For day-to-day use, this is similar to using a normal, hard-disk-based, Linux distribution.

The currently recommended method for installing Porteus and making it persistent between boots is to install from one flash drive on to another, so we did. From the user's perspective, if something is altered in the **home** directory, the change is preserved between boots. Porteus has its own system for adding applications that requires some manual file copying.

The AntiX docs recommended that we used the official installer, but we couldn't get that to work. We followed



MKUSB is a tool to create persistent USB live installations. It's easy to use, and some of these distributions recommend it in their documentation.

suggestion number two and installed from a live flash drive on to another. Our first attempt failed because the installer declared the media to be corrupted on multiple download attempts. Eventually, we were successful installing like this with a different edition of the distro. From then on, your changes are permanent.

VERDICT			
LINUX LITE	8/10	PORTEUS	7/10
EASYOS 9	9/10	UBUNTU	7/10
ANTIX	6/10		

A quirky Linux distro it may be, but EasyOS has the easiest system for installation and file and customisation persistence.

Documentation and support

How you get the distro installed and how you use it once done.

AntiX is closely connected to its base distro MX Linux and some of the resources are shared. The AntiX site has a forum with lots of sections with reasonable activity levels. However, much of the AntiX documentation is spread around between the FAQ, the wiki and other resources, and much of it is out of date and has broken web links. The documentation for the live features is still useful, once you find it.

Like most of EasyOS, the main website has a slightly scruffy, homegrown look. However, the documentation is complete, covers most areas of installation and usage, and is kept up to date by the developer. EasyOS is derived from Puppy Linux, and some solutions for that OS work for EasyOS. EasyOS has its own reasonably active forum on the Puppy Linux website.

Ubuntu is one of the best-known and most well-supported Linux distros. It has forums, tutorials and documentation on the official site and around the net. Ubuntu on a flash drive is less common, but the documentation for handling it is on the official site. Most of the time, you'd use a combination of the Ubuntu and **MKUSB** documentation if you run into problems.

The Linux Lite documentation on the website seems to be written with Linux newbies in mind. On the negative side, there is a lot of advertising, and much of the documentation feels like a



It's difficult to compete with the support that Ubuntu Linux enjoys. That said, the specific setup of Ubuntu on a flash drive is less common and less well documented.

work in progress. If you are familiar with Linux, you might be better off falling back on to Ubuntu and Debian documentation to solve problems. There is a forum with a lot of subforums and it is reasonably active. We like the first-run dialog as it contains useful links to some of these resources, online and off.

The Porteus documentation consists of a long FAQ and a wiki containing various tutorials, and we appreciate the absence of outdated material and cruft, such as general Linux information.

VERDICT			
LINUX LITE	7/10	PORTEUS	8/10
EASYOS 9	8/10	UBUNTU	9/10
ANTIX	7/10		

Ubuntu has a huge amount of support online, but it's not specific to live use, unlike the EasyOS documentation.

The verdict

EasyOS

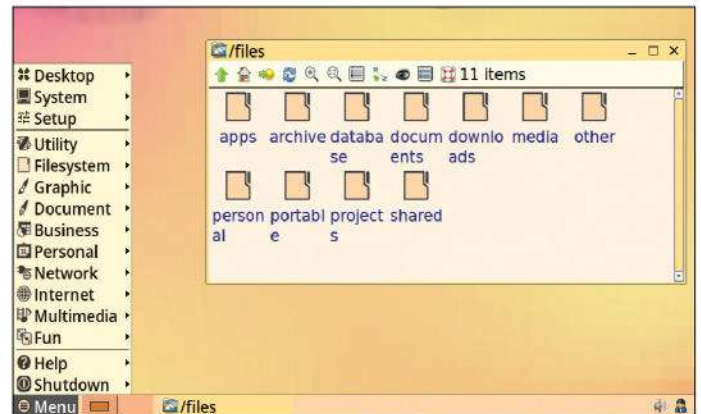
All of these distributions are useful when used as a live system, but we decided to make EasyOS the overall winner. It's a weird and wonderful distribution, but using it in a live boot context feels like its natural habitat. Nearly every part of it, from the package management to the user interface, is a little unusual, but the tools are there to work through most issues. As there are quite a few rough edges and unusual tools, it wouldn't be our first recommendation for a Linux newbie. We do like EasyOS, though. It's not a rescue distro, but it could provide you with a working computer in many difficult scenarios.

AntiX is a good all-rounder that combines a light desktop with an extensive application loadout. Because it's closely connected to MX Linux, it's far from an orphan in the Linux world when it comes to support. We followed the instructions, however, and we were a little disappointed that it took a couple of different attempts to successfully install a live, flash drive version.

Linux Lite is a distro that aims to be appealing to people who are used to Microsoft Windows as their operating system, and it succeeds in recreating a familiar experience without in any way looking like a clone of the Windows desktop. In terms of the application selection, we might have liked a few extra bells and whistles.

Porteus gives us the impression that it is designed to give you the opportunity to design the system according to your particular needs. For this reason, the software installation facilities aren't especially slick, and the customisation tools aren't standardised. Everything worked as it is supposed to when we followed the instructions, but it wouldn't be our first choice when it comes to providing the typical Linux desktop experience on a flash drive. There are also better starting points if you need a fast option that you can use to provide a Linux desktop that can handle most tasks quickly in an emergency.

Ubuntu when used in a live medium configuration isn't a setup that is carefully honed for live usage. However, the tools that are available on the live installer disk do cover most of the basic tasks, such as running office apps or playing media. It was certainly interesting to examine using it in this way, and some will be attracted to the benefits of fairly standard Ubuntu on a flash drive.



1st **EasyOS** **8/10**

Web: <https://easyos.org> **Licence:** Various

Version: 5.5.5

Bit unusual in places, but can handle most things you'd throw at a live distro.

2nd **AntiX** **7/10**

Web: <https://antixlinux.com> **Licence:** Various

Version: AntiX-23 (64-bit, Runniti)

Good all-rounder for building a normal Linux desktop on a flash drive..

3rd **Linux Lite** **7/10**

Web: www.linuxliteos.com **Licence:** Various

Version: 6.6

Can be used by anyone, but it's aimed at those switching from Windows.

4th **Porteus** **7/10**

Web: www.porteus.org **Licence:** Various

Version: 5.01 (64-bit, Cinnamon)

A bit unwieldy when used as a normal flash distro until you expand it.

5th **Ubuntu** **6/10**

Web: <https://ubuntu.com> **Licence:** Various

Version: 23.10.1

Usable live distro, even though it's not primarily designed for it.

» ALSO CONSIDER

To get even closer to a regular distro, you could install Ubuntu on a second flash drive. This might throw up performance problems and compatibility might not be quite as good if the installer has optimised the drivers for the installation machine.

Peppermint OS reminded us of Porteus in that it's lightweight, with only minimal applications, and it's ready for customisation and expansion. Although it doesn't even come with a browser initially, it does come with a built-in SSB (site-

specific browser) facility for turning websites into applications that are integrated into your desktop.

Puppy Linux is an old favourite of ours, but it hasn't been updated in quite a long time, and it seems to have been superseded by EasyOS now. There are also quite a lot of specialist distros that run from a live medium, such as security and rescue distros, but these tend to be rather compromised when used for general-purpose tasks. **LXF**

LEARN LINUX!

Finally had enough of your Windows PC? **Nick Peers** reveals how to fully road-test – then switch to – a user-friendly Linux distro.

There's a myriad reasons for why you might have fallen out of love with your computer. Perhaps you're tired of the constant updates and upgrades that slowly sap the life out of it unless you invest in expensive hardware upgrades (or even a completely new computer). Maybe it's the slow erosion of your privacy as personal information is leaked, and your computer is targeted with adverts. Or you might simply hanker after a new experience, one that

takes you back to a time when you still enjoyed computing.

The answer to all these concerns is Linux. It's supported for much longer than many OSes – up to 10 years, depending on Long Term Support options – freeing you from that constant upgrade/update cycle. Linux is also optimised for older hardware, so can give your existing PC a new lease of life. No personal data is collected, and the OS doesn't try to sneak any ads in.

Ah, you say, but Linux is difficult to use, right? Wrong. There's a multitude of Linux

distros out there written with the normal, everyday user in mind, and in this feature, we're going to introduce you to one of the most popular: Elementary OS. Its slick user interface is easy to master, and it'll instantly feel familiar. We'll show you how to test it thoroughly before installing it on your PC, either alongside your existing Windows installation or in place of it.

So, are you ready for a fresh start? Then buckle up as we take you from a Linux Learner to Linux User in around 60 minutes*. Let's go!



* Hey, we don't know how many cups of tea you'll be drinking or your internet download speeds...

Prep for a Linux OS

An ounce of preparation is worth a pound of cure.

The good news is that you don't need to commit to a full-blown install of Linux to determine if it's right for you. Many Linux distros – including Elementary OS – come with a live version of the operating system that can be run directly from DVD or USB flash drive. This enables you to get started with Linux in minutes without committing any hard drive space. Reboot your PC with the drive inserted, then look for an option to bring up the boot menu just after your PC switches on (typically F11 or F12). Press this, then select your boot media from the menu – if there's a choice, try UEFI first.

If you've installed Ventoy, you'll see a list of ISO files stored on the drive – choose the Elementary OS ISO option and it should then display a menu inviting you to Try Or Install Elementary OS. Select this to be taken through a quick language/keyboard setup wizard. From here, select Try Demo Mode and click Next.

Once Elementary OS has loaded, you can browse its desktop as if you'd installed it on to your hard drive. Sadly, you can't install any apps, but you can get a feel for how the desktop works. If you run into problems with the display, turn to page 38 for some tips.

If Elementary OS doesn't float your boat, turn to page 39 to find details of five alternative distros you can try. Download the ISO, add it to your Ventoy-enabled flash drive and give it a run. When you're ready to boot your drive, turn the page to find out how to get started with Linux!

The virtual path...

If you're not convinced by Elementary OS after running it in a live environment, the next logical step is to install it inside a virtual machine (VM) within your OS, so you can get a more realistic feel for it – you'll be able to install other software, for example.

Download and install *VirtualBox* (www.virtualbox.org), launch it and choose Machine > New. Click the



down arrow next to ISO Image to select your Elementary ISO file and give your virtual machine a suitable name. Expand Hardware to allocate more memory and CPU resources if available – a couple of cores and 4GB of memory is ideal. For the Hard Disk increase the size of the virtual hard disk to a more comfortable 64GB, plus choose a suitable drive or folder on your PC to house it all.

Click Finish, then click Start and watch as Elementary runs inside its own window. When you come to the Try Or Install screen, choose Erase Disk And Install, and follow the guide in the main copy to install a fully working version inside the virtual machine (ignore the warning about it not working properly in a VM by clicking Install Anyway).

Once installed, shut down the VM and click Settings > Display. Verify VMSVGA is selected as the Graphics Controller, tick Enable 3D Acceleration and allocate the maximum amount of RAM. Boot into your VM and follow the optional steps outlined at <https://i12bretro.github.io/tutorials/O137.html> to install the Guest Additions and improve performance.

Elementary OS can be spun up in a virtual machine in seconds, so you can test run Linux for yourself on almost any system.

» BUILD YOUR INSTALL MEDIA

Like all Linux distros, Elementary OS's installation media is supplied as a single ISO file, which you can burn to DVD or flash to a suitable USB thumb drive (4GB or larger) using a tool such as *Etcher* (<https://etcher.balena.io>). But if you have a larger flash drive, why not use a special tool called *Ventoy* (<https://ventoy.net>) to create a single flash drive

on which to house all your bootable media?

On your Windows PC, download the latest build from the website and use it to set up a single large flash drive (16GB or bigger) to house all your bootable media. Once created, all you need to do is copy your downloaded ISO files to the drive. When you later boot your PC from the flash drive,

Ventoy displays a menu letting you choose which ISO to boot from, making it easy to switch between them.

Now return to <https://elementary.io> – type a zero into the Custom box and click Download Elementary OS. Save the file to your hard drive, then copy it across to your *Ventoy* boot drive. After creating your fail-safe backup using *Hasleo Backup*

Suite (see the main text), copy the ISO file containing your emergency bootable media to the flash drive, too.

If you'd like to try out any other Linux distributions (see page 39) and your thumb drive has enough free space left, you can download the required ISO file from each website and copy one or more to your thumb drive, ready for trying out.

Take the Elementary tour

Discover how to get started using your new flavour of Linux.



Booting into Elementary OS – using Live Distro, a virtual machine (see page 36) or full install – you should find yourself at the login screen, where you're prompted to enter the password you just created to log into your new OS for the first time. When the desktop appears, you'll see a verdant mountain landscape as your wallpaper, highlighting how Elementary is designed very much in the mould of Mac OS.

At the bottom is the Dock – it's a more elegant form of Windows' taskbar, while there's a Panel at the top. On the left of this, click Applications to open Elementary's launcher, its equivalent of the Start menu. You'll see a series of shortcut icons spread over two pages (roll the mouse wheel up and down to move between them) – clicking one opens it, while right-clicking reveals options to pin or remove it from the Dock, uninstall the app or view it in the AppCenter.

The AppCenter is Elementary's app store, similar to the Microsoft Store in Windows or your phone's app store. It's not the only way to install programs on your PC – and is, in fact, quite lightly populated with apps written specially for Elementary's Pantheon desktop. We'll cover both it and alternative app sources shortly.

Continuing our tour, in the middle of the Panel is the date and time – roll your mouse over this to reveal a pop-up monthly calendar with a list of any events you've set up using Elementary's main *Calendar* app or via any online CalDAV service you've connected to it.

The top-right corner of the Panel reveals a series of buttons, each of which reveals a pop-up menu. From left to right: a speaker button lets you control audio



Elementary's desktop is called Pantheon and was designed with Mac OS switchers in mind. It's bright, colourful and easy to use.

volume and music playback via Elementary's *Music* app; a network button makes it easy to quickly enable and disable networks; a notifications pop-up offers a 'do not disturb' option; and a power button make it easy to lock, log out, suspend, shut down or restart your system.

All four pop-ups also contain shortcuts to their relevant entries in the System Settings menu, which can also be accessed from the launcher and the Dock.

Set up your system

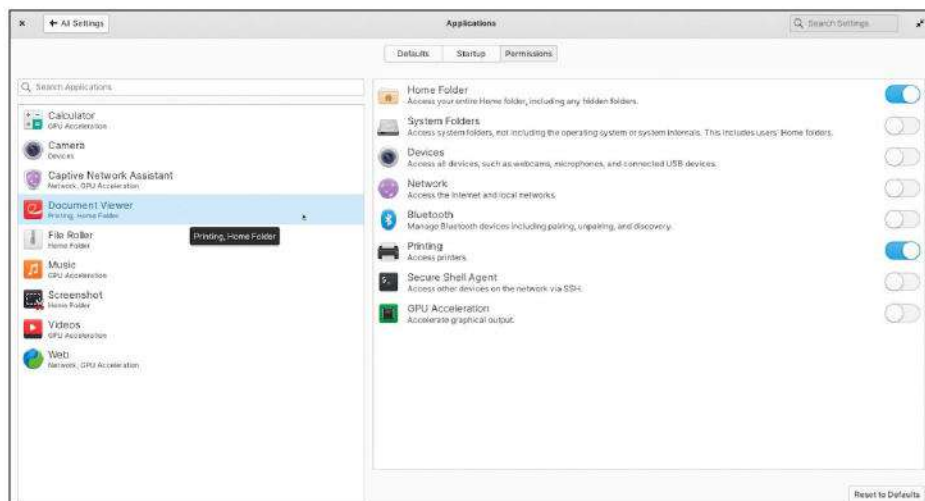
Take the time to go through System Settings – this is where you get to fine-tune Elementary to your needs. It's a little spartan compared to other OSes – there are 21 sections split into four categories: Personal, Hardware, Network & Wireless, and Administration.

For the most part, things are self-explanatory – Applications is where you can switch defaults to new programs if you replace them. It's also the place to go to manage startup entries and set permissions for individual apps, managing their access to certain folders and devices, such as webcams and printers.

If the layout isn't quite to your liking, explore the options under Desktop, from switching wallpaper and accent colours to resizing the Dock and even improving readability for those with dyslexia.

Visit Displays to resize the screen if necessary or enable night light support and apply filters to make the screen easier to read and less distracting. Go to Online Accounts to connect CalDAV calendars and tasks as well as IMAP email accounts to use with the built-in *Mail* and *Calendar* apps. A Sharing section enables you to switch on DLNA servers for streaming music, photos and videos to compatible devices.

The System Settings tool is worth exploring in depth after installing Elementary – it works in a similar way to Windows' equivalent.



Another unique feature can be found under Screen Time & Limits, where if you share your PC with others, you can apply screen time limits, plus block selected applications and websites – perfect for children.

Some sections are greyed out with a warning that administrator rights are required – click the Unlock button and enter your user password to make changes.

Find more apps

While Elementary comes with a core set of apps similar to those bundled in Windows, you'll soon be itching to do more with your new OS. You can install apps several ways in Elementary – just like you can in Windows – but the simplest place to start is its AppCenter, accessible directly from the Dock.

Click this to open an app store not dissimilar from that on your phone. There's a search tool at the top, a list of featured apps, and categories for you to browse. Click one, and a list of paid-for and free apps is shown. You'll also see an orange flash button at the top – click this to see a list of all installed apps (these include the apps preinstalled with Elementary) as well as any that have updates available. Click Update next to individual entries, or Update All to update them all in one go, or click the Settings button next to the flash button to reveal an automatic updates option.

You may be disappointed with the slim pickings on offer – try a search for a well-known tool such as *LibreOffice* and you're likely to come up short, although you are directed to an alternative app centre, Flathub, through Elementary's web browser. Flathub has many more apps to choose from – click Install to download a small file with a flatpakref extension. Open your **Downloads** folder in the *Files* app, then double-click this. A dialog warning you about this 'untrusted' app pops up, with a summary, download size and warning that updates aren't checked before they're offered. Tick I Understand and click Install Anyway to proceed.

Apps installed through Flathub also appear in the AppCenter alongside other installed apps, but there's one major disadvantage to Flathub: because of the way apps are packaged (to run in a

sandbox for security and compatibility purposes), they tend to be much larger than apps from elsewhere.

Thankfully, apps can be added from other sources, too – because Elementary is effectively a variant of the massively popular Ubuntu distro, it supports Ubuntu repositories, which are collections of apps hosted in a single location. These can be installed directly through the terminal using `sudo apt install appname`, but if you don't know the application name, a simpler option is to install another package manager alongside AppCenter. Open the terminal and issue the following commands:

```
$ sudo apt update
```


```
$ sudo apt install synaptic
```

Once installed, close the terminal and open the Applications launcher, where you'll see an entry for *Synaptic Package Manager*. Open this and you'll gain access to a much wider range of programs.

You can also source some programs from their own websites. Look for two types of file: DEBs are packages similar to Windows installers, while AppImages are like Windows portable apps. Search *Synaptic Package Manager* for a program called *GDebi*, which simplifies the install process for DEB packages.

When it comes to AppImages, the files need to be tweaked to make them executable in Linux. To do this, right-click the file in *Files* and choose Properties > Permissions before clicking Execute under Owner, which changes the permissions to 764. Click Close and the program should now run when you double-click it.

The AppCenter isn't just for finding and installing new programs; use it to manage and update existing ones, too.



USING THE TERMINAL

While Elementary OS provides an easy-to-digest desktop to explore and use, there's no getting away from the fact that it's a Linux operating system, and that means there's an underlying command-line interface (shell) available. You interact with it via the terminal, accessible via the launcher.

The shell uses a language called Bash, and the basics

are quite easy to grasp as its syntax is relatively simple: `utility command -option`. But rather than try to cram a crash course in here, go to www.terminaltutor.com, where Niklas Wenzel has built a brilliant site that steps you through the basics with an interactive shell you can type commands into.

When you open Elementary's terminal, you'll

see it's a relatively simple affair. You open a session to the shell pointing to your home (~) directory, logged in as your own user. You can open additional sessions using the + button on the left – each one runs in its own separate tab. A Search button enables you to search each tab for keywords, while the Settings button provides a handy tool for changing the

size of the text as well as the terminal's colour scheme.

You may feel like the terminal is something you'd rather avoid, but it's worth exploring it a little bit, simply because it's often much quicker to type (or paste) commands into the terminal to achieve things instead of pointing and clicking your way through various application interfaces.



Install Elementary OS

Discover how to safely install Linux on your PC.

If you've liked what you've seen and decided to install Elementary OS, it's time to take the plunge – or, if you're not quite ready, check out the box (below) on installing Elementary in a virtual machine, which has the added benefit of giving you a dry run through the install process without having to worry about damaging your system.

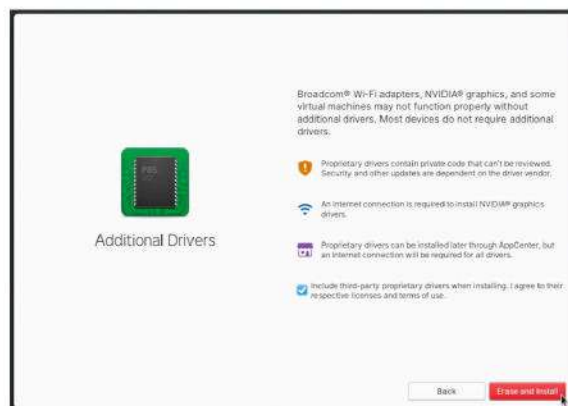
Now consider where and how you're installing Elementary OS. In most cases, you'll be installing it alongside Windows. This can be done one of two ways: either install Elementary on its own separate hard drive, or partition your existing system drive to make space for Elementary alongside Windows.

If you choose the latter option, you need to shrink your Windows drive to free up a minimum of 64GB of space for Linux. Make sure you leave Windows some free space, too – we recommend your current drive has a minimum of 100GB of free space before attempting to partition it. One prerequisite of a dual-boot setup is that you disable hibernation in Windows, which has the bonus of freeing up gigabytes of drive space (the equivalent of however much RAM is in your PC). To do this, right-click the Start button and choose Terminal (Admin). Type the following command and hit Enter:

```
$ powercfg -h off
```

After you've cleared out enough space, right-click the Start button and choose Disk Management. Right-click your Windows drive and choose Shrink Volume to attempt to shrink it by 64GB or more. If this isn't possible due to what Windows terms "unmovable" files, use a free third-party tool such as *Paragon Partition Manager Community Edition* (www.paragon-software.com/free/pm-express) to free up the space.

If your drive isn't big enough to accommodate Linux as well as Windows, and there isn't space to fit a second drive, you may need to consider upgrading to a larger one – if you go down this route, use *Hasleo Backup Suite's Clone > Disk Clone* option to migrate



Tick the box when prompted to install proprietary drivers that will simplify hardware setup on many systems.

your existing Windows installation to the new, larger drive before proceeding.

Start the install process

You're now ready to install Elementary. With your USB drive plugged in, restart your PC. The start of the process is identical to giving Elementary a test drive, so repeat the process of booting from your USB drive from the previous page.

You'll now find yourself back at the Try Or Install screen. At this point, you have two broad choices: Erase Disk And Install, or Custom Install (Advanced). The first option is the one to choose if you're either wiping Windows completely from your system or installing Elementary to its own drive. This latter option is our preferred method – not only do you get to preserve your Windows installation, but by keeping Elementary on a separate drive, you don't have to get involved with any partitioning.

But this isn't always practical – for example, you want Elementary alongside Windows on a laptop with

» BEFORE YOU INSTALL...

Even if you plan to install Elementary alongside Windows, there's a chance the partitioning might fail, so it pays to take a fail-safe backup before you begin. This just makes for good digital safety. The best free tool for the job is *Hasleo Backup Suite Free* (www.easyuefi.com). Download and install it, then launch the app. Create your emergency boot

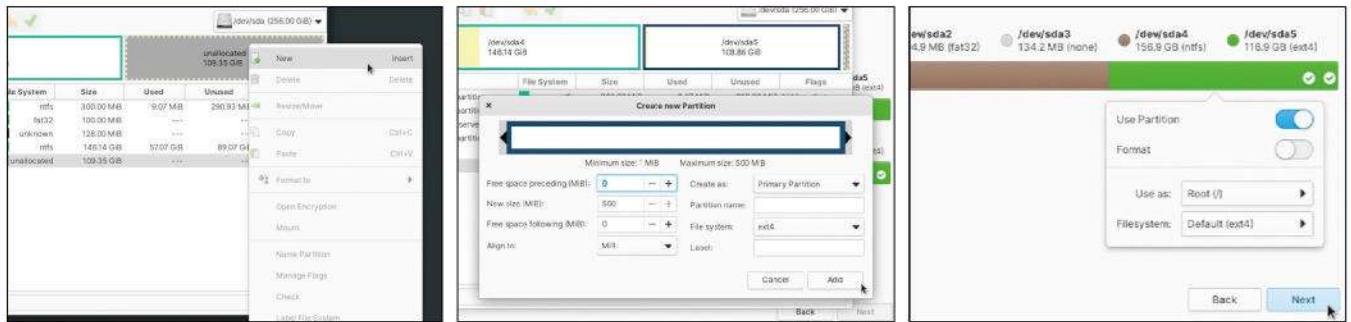
media – navigate to Tools > Emergency Disk. Tick Download WinPE Components, then click Next and wait for the rescue disc to be created. Leave Export ISO selected and choose where to save the file. Once done, see the Build Your Install Media box (below) for a way to copy it and your Linux installation media to the same drive.

Switch to Backup, where you'll see three options; start with System Backup. This preselects all disks and partitions required to run Windows. Choose a suitable backup location (a different physical drive – preferably external – to the one Windows is installed on) and click Proceed. Once the backup has been taken, go back to Tools and select

Check Image to verify the backup can be relied on if you ever need it.

Now ask yourself if you need to back up anything else – if your data is stored on a separate partition or drive from Windows itself, it's worth either taking a full Disk/Partition Backup of everything on that drive, or choosing File Backup to back up selected files or folders.

PARTITION FOR A DUAL-BOOT SETUP



1 Select unallocated space

After clicking Custom Install (Advanced), you're shown a graphical representation of your hard drive's current partition structure – the unused part is what you need to set up. Click Modify Partitions to open the *GParted* partition editor. Right-click the unallocated space as shown above and choose New.

2 Create required partition

Type 500 into the Free Space Following (MiB) box to allocate all but the last 500MB of the drive to Elementary. Leave the other settings as they are (Primary Partition, ext4 filesystem) and click Add. Now right-click the unallocated space and choose New followed by Add. Choose Edit > Apply All Operations.

3 Assign partition

Click Apply, then Close when finished. Choose *GParted* > Quit. Right-click the small partition on the right and flick Use Partition to on. Set Use As to Boot (Boot/EFI). Now right-click the partition to its left and flick that switch on, too. Leave it set to Root (/) and Default (ext4), then click Next.

only one physical hard disk, which is where the advanced option comes in. This lets you configure the free space to welcome Elementary on to your PC.

Installation proper

The process diverges depending on which install option you choose. If you select Erase Disk And Install, you're shown a list of all detected drives. Whichever drive you select is wiped clean and Elementary is installed on it, so choose with care. If installing Elementary to its own drive, use the drive size to identify the correct one to erase before clicking Next.

If you're choosing the Custom Install (Advanced) option, follow the step-by-step guide (above) to see how to carve out the freed-up space for Elementary.

Once your drive has been selected or partitioned, you may be prompted to enable drive encryption for security purposes – either select Choose Password to do so (but make sure the password is both strong and memorable – store it somewhere safe, such as your password manager) or choose Don't Encrypt if you don't need it.

Next, you're prompted to install proprietary drivers – we recommend ticking the box before clicking Erase And Install to begin the installation proper. Wait for


drives to be erased, files copied and extracted, and the installation to proceed.

Post-install steps

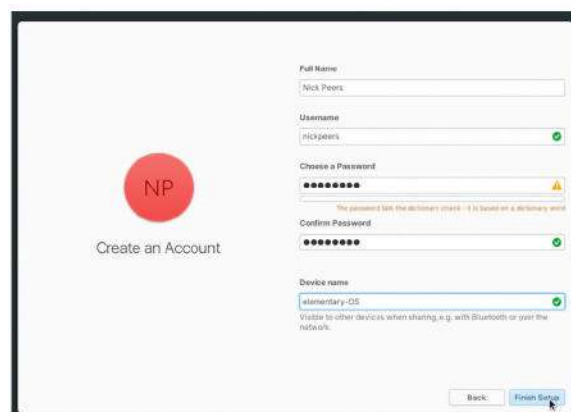
You're next prompted to unplug your flash drive (this is done automatically if you're installing in a virtual machine) and press Enter to reboot. Do so and if you're lucky, you're shown the GRUB boot menu for the first time, where you're given a choice between booting Elementary or Windows every time you start your PC.

If this doesn't appear and you find yourself back in Windows, restart your PC and press the key prompted (typically Del or F2) to enter your system BIOS/UEFI. Locate the Boot section, where you should see that your Windows boot manager is currently the default boot device – change this to the Elementary OS entry, save your changes and exit. You should now boot to GRUB – leave Elementary selected to complete setup.

This can take some time in a VM (wait if you find yourself halted at a terminal login prompt, for example). Eventually, the wizard reappears – skip through the language settings, then set up your user account (full name, username, password, and computer name for identifying on the network), then click Finish Setup. Congratulations you've fully installed Elementary OS.



After installing, it's likely Windows will reassert itself as the default choice at startup. Change this in the UEFI.



Apart from restating your language and keyboard choices, you need to set up your own user account post-install.



Solve hardware issues

Fix classic problems first-timers have with displays, networks and more.

Hopefully your introduction to Elementary will be a smooth one, but if you run into problems when running on real hardware don't panic. As with Windows, most issues are down to drivers, but if you tick the option to install proprietary drivers during setup, you should avoid most issues.

One issue you can't avoid if you're trying to run the Elementary OS live environment on a PC with Nvidia graphics is the fact that it boots to a virtually unreadable desktop that's bright yellow with white text. The solution is to reboot, then choose Try Or Install Elementary OS (Safe Graphics) at the first boot menu.

If you subsequently don't tick the proprietary driver box when installing Elementary, you're faced with the same virtually unreadable desktop. To fix this post-install, add the required Nvidia drivers from a terminal window. Enter the following two commands, then reboot:

```
$ sudo apt install linux-headers-$(uname -r)
$ sudo apt-get install nvidia-driver-525
```

If you suffer from other less severe colour issues, the likely culprit is a faulty

monitor colour profile. You can try navigating to System Settings > Displays > Filters and experimenting with the options there for a temporary fix, but a more permanent solution involves using a terminal-only app called *Colormgr* with the following commands:

```
$ colormgr get-devices-by-kind display | grep "Device ID"
$ colormgr find-profile-by-filename '/usr/share/color/icc/colord/sRGB.icc' | grep "Profile ID"
```

In the following line, substitute Device ID and Profile ID with the values returned by the first two commands:

```
$ colormgr device-add-profile YourDeviceID
YourProfileID
```

If this solves the issue, make the switch permanent:

```
$ colormgr device-make-profile-default
'YourDeviceID' 'YourProfileID'
```

Fix networking issues

Another bugbear is getting your PC connected to the network and thus the internet. Driver issues should only affect certain wireless adaptors, and hopefully they should be covered by the proprietary drivers, but sometimes you need to jump through more hoops.

Click the Network button in the top-right, then click Network Settings. This should verify whether your wired or wireless network adaptor has been detected (it's in the list if it is) and if it's connected. In most cases, the automatic settings should get you connected, but if you need to configure the connection manually, click Advanced Settings. For Ethernet connections, focus on the IPv4 Settings tab for setting IP addresses.

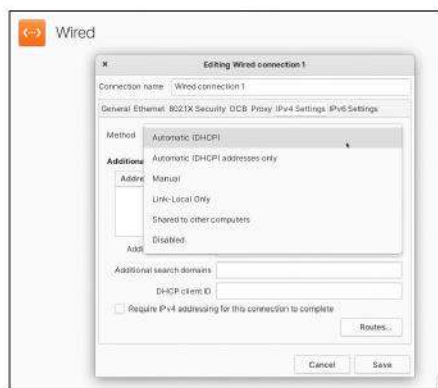
If your adaptor isn't present, you need to track down drivers. If you don't know its make and model, input one of the following commands, depending on its type:

```
$ lspci -v | grep Ethernet
$ lspci -v | grep Wireless
```

Visit the manufacturer's website for info or search the web for Ubuntu 22.04 installation instructions (the procedure should be the same in Elementary).

Printer issues

If you're unable to get your printer working via System Settings > Printers, see our guide to working with printers in **LXF309**, or visit <https://openprinting.github.io> for a guide to the new OpenPrinting protocols.



If your network adapter is detected, but you still can't connect, visit the Network section to set it up properly.

TROUBLESHOOT PROBLEMS ONLINE

If you run into problems with Elementary OS, you are unlikely to be alone. Thankfully, there are plenty of online resources to help you track down a solution. A general web search is usually the best place to start – if there's a specific error message, start by searching for that with the words 'elementary' and 'os', or provide keywords providing a more general description (such as 'display yellow

elementary os'). Most results point to user forums, which are worth bookmarking for future reference.

Elementary's own support channels at <https://elementary.io/support> have a list of apps and topics that redirect you to the relevant part of the specific app's GitHub page. If nothing appears under Q&A, check the Issues section. Most hardware-related queries can't be answered here, so

click Other Questions & Help to be redirected to a more general Q&A section.

Other places worth bookmarking include channels on StackExchange (<https://elementaryos.stackexchange.com>) and Reddit (www.reddit.com/r/elementaryos). If you can't find a solution across any of these, consider signing up with one or more sites to post your query. Failing that, try filing an issue through the

Feedback app (see <https://docs.elementary.io/contributor-guide/feedback/reporting-issues>) or emailing our resident expert at answers@linuxformat.com to see if he can help.

Remember to include as much relevant information as you can – your hardware specs for example, particularly the make and model of any troublesome hardware like your graphics card or network adaptor.



Try some distro hopping

Don't like Elementary? Try one of these beginner-friendly OSes.

While it's a potentially easy way into Linux, you may decide that Elementary OS isn't for you – perhaps you don't like its desktop and cut-down app store or find it's more trouble getting set up than it's worth. But this doesn't have to mean bringing your adventures into the world of Linux to a premature end. The great thing about Linux – unlike Mac OS or Windows – is that you have an unbelievably diverse choice of alternative distros to choose from.

Most distros are based on another distro, tweaked to provide a specific user experience, or fill a perceived gap in the market. Most – including Elementary OS – are based on Ubuntu, one of Linux's most popular distros (itself based on another distro – Debian). Ubuntu is backed by Canonical, a well-established company capable of providing lots of resources and support. And if there's one area of the Linux marketplace where there's plenty of choice, it's those distros looking to appeal to newcomers to Linux.

We've rounded up five alternative distros to check out – the only cost to

you is time downloading and testing each one. All five are distributed as ISO files offering a live environment and an installer, so just download them to your hard drive, then either create separate installation media for each or copy the ISOs to your Ventoy USB flash drive and switch between them to see which take your fancy.

You'll see that many distros offer multiple options, each sporting a different desktop. There are desktops like Elementary's Prometheus, offering a slick, modern experience, as well as desktops designed for lower-end machines, like the minimalist Xfce. You'll also see that desktops work across multiple distros, enabling you to mix and match as you see fit. Happy testing! **LV**



Linux Mint (<https://linuxmint.com>)




Pop!_OS (<https://pop.system76.com>)



Ubuntu (<https://ubuntu.com>)



Zorin OS (<https://zorin.com/os>)



Manjaro Linux (<https://manjaro.org>)

FIVE ALTERNATIVE DISTROS TO TRY

Linux Mint

<https://linuxmint.com>

For years, Linux Mint was considered the best distro for Windows switchers, thanks to its desktop, which mirrored Windows 7. While Elementary is slicker, Linux Mint is robust and well supported, and its Cinnamon edition should be your first port of call if you fail to take to Elementary.

Pop!_OS

<https://pop.system76.com>

This distro comes from Linux PC manufacturer System 76 and is tailored to STEM and creative professionals. However, it's still a good choice for Windows switchers and anyone who wants to prioritise using their computer over spending ages configuring, tweaking and ultimately troubleshooting it.

Ubuntu

<https://ubuntu.com>

Like Elementary itself, most of the distros featured here are based on Ubuntu, so why not give the original a try when testing different flavours of Linux? The Gnome desktop will be different from what you're used to, but doesn't take long to master, and Ubuntu is the best-supported OS of any in this roundup.

Zorin OS

<https://zorin.com/os>

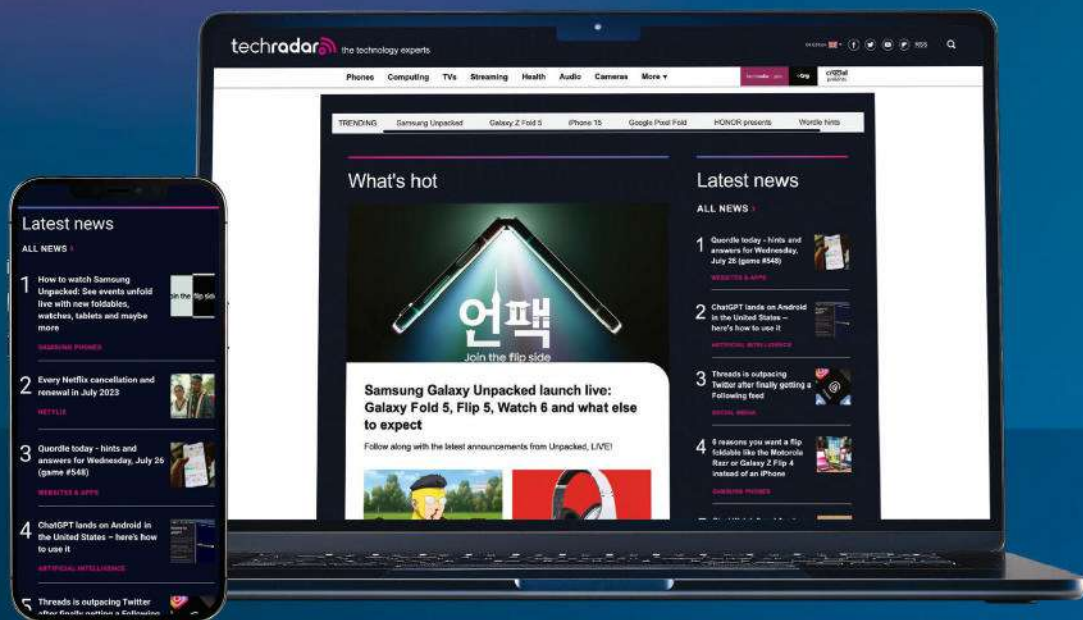
Zorin OS is aimed at switchers, to the point of highlighting the fact that its *Appearance* app is geared towards letting you fine-tune its desktop to resemble the version of Windows or Mac OS you're migrating from. Even so, its Ubuntu origins remain unmistakable behind the slick desktop interface.

Manjaro Linux

<https://manjaro.org>

Unlike the other distros we've featured, Manjaro is based on Arch Linux. While Arch Linux is considered difficult to learn, it's fast and powerful, and Manjaro is a specially built variant with user-friendliness in mind. It comes with a choice of desktop, including the slick and modern Plasma desktop from KDE.


Meet the technology experts



- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

techradar
the technology experts

www.techradar.com



Arm buys a slice of Pi

The deal effectively cements the Raspberry Pi to the Arm ISA for an undisclosed sum.

Arm Holdings has announced that it has made a strategic investment: a minority stake in Raspberry Pi.

The deal extends the long-term partnership between Arm and Raspberry Pi, which has seen Arm CPUs feature in all of the Pi and Pi Pico SoCs. The partnership began way before the Pi was available for sale, in 2008, when the original board was still just a dream. Fast-forward to 2023 and we have a generation of learners who have taken their first steps with coding, science and electronics thanks to the Raspberry Pi.

"Arm and Raspberry Pi share a vision to make computing accessible for all, by

lowering barriers to innovation so that anyone, anywhere can learn, experience and create new IoT solutions," said Paul Williamson, SVP and GM, Internet of Things Line of Business, Arm.

Arm's minority stake in Raspberry Pi also shows a firm commitment to the continuation of Arm CPUs in future Raspberry Pis. With the rise of RISC-V CPUs in devices ranging from \$9 to hundreds of dollars, it is clear that we will not be seeing a RISC-V-based Raspberry Pi in the foreseeable future.

See the official release: <https://newsroom.arm.com/news/raspberry-pi-investment>.



Les Pounder works with groups such as the Raspberry Pi Foundation to help boost people's maker skills.

» I'LL NEVER GET MY FILL OF PI

The Raspberry Pi 5 was announced on 28th September, but it wasn't on general sale until 23rd October. Well, it was if you pre-ordered. It seems that the Raspberry Pi 5 is the victim of its own success and stocks are thin. Pre-orders are being fulfilled and Raspberry Pi CEO Eben Upton has provided an update to say that the Sony facility in South Wales is working hard to keep up with demand. If you pre-ordered on the day, nay the hour of release, then chances are that you already have or will soon receive a unit. If you left it for a few hours, then you'll have to wait until the end of November or early December. If you waited even longer, well, 2024 isn't too far away now...

So, was the Raspberry Pi 5 worth the wait? Of course. For not much more money than a Raspberry Pi 4, we get a far more powerful CPU and GPU, along with the much-anticipated PCIe breakout. This breakout alone is worth investigating. Soon the market will be flooded with M.2 HATs and PCIe Ethernet, and someone (I'm looking at YouTuber Jeff Geerling) will find a way to connect a GPU to the Pi 5.

Is the 'elderly' Raspberry Pi 4 consigned to the trash? No way! The Raspberry Pi 4 has plenty of processing power and it is compatible with all of your HATs, unlike the Raspberry Pi 5, which introduced the RP1 chip and Bookworm-based OS, which makes using HATs a sore subject. For now, I'll be using my Raspberry Pi 4 for many years to come. A heavily modified Pi 4 powers my arcade cabinet.




CREDIT: Jon Thackray CC BY-SA 4.0, https://commons.wikimedia.org/wiki/File:Arm_ABCD_building.jpg

NASA stuff

Super-pressure Pi.

A NASA team has been using a Pi 3B to help collect and back up data for its SuperBIT telescope. The balloon-launched device busily gathered data from around our home planet for a month. The original plan was to use the SpaceX Starlink system, but after failing, they implemented a backup delivery method using the Pi. Get the science at: www.mdpi.com/2226-4310/10/11/960.



CREDIT: NASA

Useful Micro:bit

Dig them all out!

In an attempt to resurrect the success of the BBC Microcomputer, the BBC put its weight behind the release of the Micro:bit back in 2016. To help keep these devices alive, the Pi Foundation has a raft of new tutorial material, so go check it out: <https://bit.ly/lxf310micro>.



It's still useful!

Pi OS Bookworm

Les Pounder has a Raspberry Pi 5 and isn't afraid to use it, now that he has the latest Raspberry Pi OS.

IN BRIEF

This is the latest Raspberry Pi OS, built from Debian 12 Bookworm. We have a solid desktop experience, but issues with how Python works and the GPIO mean that makers will want to hold off until things are resolved. Now based on Wayland, Raspberry Pi OS has a fresh and fast feel to it, the most grown-up Raspberry Pi OS since its release.

The new Pi 5, our first flagship board since 2019, means we need a new OS, and Raspberry Pi OS Bookworm has hit the mirrors and is ready for download.

Raspberry Pi OS Bookworm is available for all models of Raspberry Pi, but your best experience is going to be with a Pi 4 or 5. Bookworm is offered as 32 and 64-bit images, with the latter being for Pi 3 onwards. Via the official *Raspberry Pi Imager*, we have to select No Filtering under Raspberry Pi Device to write a card for an older, 32-bit Pi.


You really need a fast microSD card (or USB 3 drive), especially with the Raspberry Pi 5's new SDR104-compatible microSD card reader. Long story short, the Pi 5 and an SDR104 microSD card is as fast as a USB 3 SSD on a Pi 4. We did a quick boot test using a class 10 Sandisk microSD and it booted in 33.19 seconds. Not bad, but an SDR104-compatible card dropped that to 23.26 seconds. An NVMe SSD via USB 3 saw 19.25 seconds.

The desktop doesn't look much different from the previous release, but we have a new Wayland-based desktop that uses Wayfire for the general UI. X11 is no more, kind of. There is a tool, *XWayland*, which intercepts any apps that require X11 and helps keep your older apps running. Best of all, it does this automatically.

For the first time, we see a Pi-optimised version of *Firefox*. We found it performed better than *Chromium* (also installed). We did a YouTube test and a 1080p video still dropped frames, but without 'stats for nerds', we couldn't tell the difference. 4K streaming video playback is still lacking but local playback, using *VLC* and the h.265 codec, offers decent results; using h.264 results in a beautiful slideshow but not a watchable video.

Audio is now handled via PipeWire and it just works. We connected our Bluetooth speaker and had music blasting across the office.

Let's address the elephant in the room: Python. Debian 12, which this release is based upon, has changed

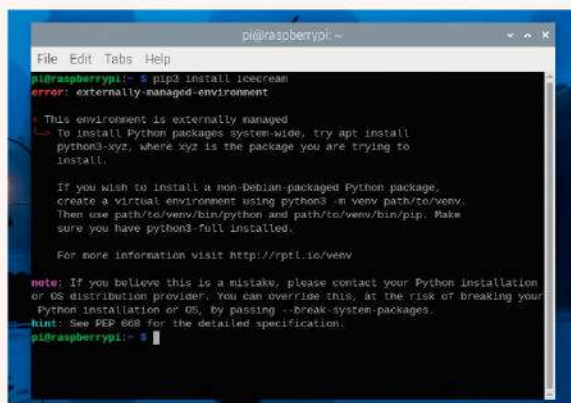


It looks great and it packs plenty of features without bamboozling us with icons. Bluetooth, networking and updates are all within easy reach.

how it handles externally managed Python modules. PEP 668 prevents the user installing Python modules at an OS level using the Python package manager *Pip*. If you use *Apt* to install a Python module, all is well. But using *Pip* triggers an error message that the Raspberry Pi developers have enhanced with a message and URL for more information. The information directs us to create Python virtual environments, containers for our Python project code that keep our code away from the OS-level Python install. This is a smart move, but for the target audience, it seems like something we could do without.

Another elephant has walked into the room and this one has GPIO written all over it. It seems that *RPi.GPIO* support is an issue. We could not get any third-party HAT to work with Bookworm. Raspberry Pi's *Sense HAT* works, but that relies on *Igmpio*, an alternative to *RPi.GPIO*. The GPIO does work – we can use *GPIO Zero* to make projects – but if you rely on HATs for your project, stick with an older release (and Pi) until everything is resolved.

Overall, Raspberry Pi OS Bookworm feels like a desktop operating system. It just works with the glaring issues of the GPIO and Python virtual environments. There is a lot to love in this release, but GPIO hackers had better hold off, for now. **LXF**



If you try to install a Python module to the OS-level Python install, PEP 668 kindly advises you of the new way of working.

VERDICT

DEVELOPER: Raspberry Pi
WEB: www.raspberrypi.com/software
LICENCE: Mixed

FEATURES	8/10	EASE OF USE	8/10
PERFORMANCE	8/10	DOCUMENTATION	8/10

It looks great, feels smooth and performs really well. Now please fix the GPIO issue and we will be happy.

» **Rating 8/10**

Creality K1 Max

A big fan of fancy, **Denise Bertacchi** has taken to the K1's bigger brother.

SPECS

Vol: 300x300x300mm

Type: PLA, PETG, TPU, ABC (up to 300°C)

Extruder: Direct drive

Nozzle: 0.4mm

Platform: Smooth PEI steel flex plate, heated

Levelling: Auto with LiDAR

Sensors: Runout, build camera

Comms: LAN, Creality Cloud, USB

Control: 4.3-inch colour touchscreen

Size: 435x462x526mm, 18kg

Creality is one-upping itself with the K1 Max, a bigger 300x300x300mm Core XY 3D printer. Like the K1, this is an enclosed single-colour printer, capable of print speeds up to 600mm/s.


The K1 Max isn't just a bigger printer, it's a more deluxe printer, with LiDAR and an AI camera to warn you if a print fails. Like the K1, it's running a limited version of *Klipper*, though Creality has recently released an open version of the firmware after facing community backlash for restricting access. If you're interested in using vanilla *Klipper*, you need to install it yourself.

The Creality K1 Max is a modern-looking, fully enclosed Core XY printer. Naysayers call it a Bambu copycat, but it's a larger machine, capable of running open source *Klipper*. Its build plate is 300x300mm, with 300mm build height, a size typical of Creality's older CR10 models. The build plate is a smooth PEI-coated sheet that requires glue stick – annoying, but it works.

It uses a flat glass lid, a carbon filter in the cabin, a camera and LiDAR. It uses those last two to assist with bed levelling and first-layer perfection. The printer uses a form of AI to scan the build plate and then the first layer to make sure everything is going to plan. It continues to monitor the print process and can warn you if it detects a print failure. This mode needs to be switched on and only works in conjunction with the Creality Cloud. If operating in LAN-only mode, you have to monitor prints the old-fashioned way – with your own eyeballs.

We're not big fans of any machine with a rear-mounted spool holder, which requires turning the printer around to access the rack and the filament runout sensor when loading materials. Interestingly enough, one of the models included with the printer is a freestanding filament rack, which we immediately printed.

The K1 Max ships with an improved hotend and tweaked extruder, and throughout our testing worked flawlessly. Like the K1, this machine has excellent cooling and an enclosure that makes printing ABS and ASA very easy. Bed levelling and input shaping are automatic and worry-free. It's not a quiet machine, but that seems to be



the price for speed. It's not the sort of machine you want to have in your bedroom or next to your TV.

The screen and menu system are easy to use. When paired with *Creality Slicer*, you get wonderful thumbnails of your prints, which make it easy to select jobs. The K1 Max remembers everything it has printed until you clear the memory, so it's great for running recurring parts.

The Creality K1 Max is 99% assembled and only needs to be uncrated – yes, uncrated – and unpacked. To protect the glass sides, Creality ships the K1 Max in a wooden crate with standard cardboard packing inside.

The Creality K1 Max has a fully auto-levelling program that includes input shaping with the help of built-in sensors. Push the button and the machine does the rest.

We ran a Benchy using Speed Benchy rules: 0.25mm layer height, two walls, three top and bottom layers, 10% infill. We turned off combing and z hop, and let her rip using the default speed of 300mm/s.

The Benchy looks better in person, but you can see that the layers are fairly even, the curves are nice, and the overhangs are near perfect. This Benchy printed in an eye-popping 17 minutes and 38 seconds, the fastest Speed Benchy we've printed. An Ender 3 S1 running at its top speed of 100mm/s would take 55 minutes. **LXF**

Destroying the world one smoothly printed Benchy at a time.



VERDICT

DEVELOPER: Creality

WEB: www.crealityofficial.co.uk

PRICE: £859

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	9/10	VALUE	7/10

The Creality K1 Max is an excellent printer for makers who want a fast, fully enclosed Core XY 3D printer that's bigger than what Bambu Lab can offer.

» **Rating 9/10**

RASPBERRY PI 5

Take your Pi 5 to the next level!

Les Pounder has got his hands on a Raspberry Pi 5 and wants to show you how to get the best from it.



OUR EXPERT

Les Pounder is associate editor at Tom's Hardware and a freelance maker for hire. He blogs about his adventures and projects at <http://bigles.com>.

Your shiny new Raspberry Pi 5 is the fastest flagship Pi ever. But can we make it a little better? Sure we can, we just need to overclock it, boot from USB and make sure that it is adequately cooled.

In this tutorial, we will overclock the Pi 5 and show you how to install the new Raspberry Pi OS to a USB 3 drive. We'll also pass on a few pearls of wisdom from our weeks of testing and using the latest Pi.

Cases and cooling

The official Raspberry Pi 5 case looks pretty much like the Raspberry Pi 4 model, except for the Ethernet port being switched to the other side of the USB ports, and the new case has built-in cooling. Yes, the Raspberry Pi 5 needs cooling. It can run without it but you will soon encounter CPU throttling. The official case has a tiny fan to blow cool air over the CPU, Wi-Fi and the new RP1 chip. It connects to the new fan connector, meaning that it has full access to the GPIO. We've got a case on order, as well as Pimoroni's new PiBow Coupe case for the Raspberry Pi 5.

For the best cooling, you will want to spend the extra £5 and pick up the Official Active Cooler. Looking resplendent in anodised aluminium, this cooler has a heatsink and active fan that connects to the CPU, Wi-Fi and RP1 chip via thermal pads. Held in place using new mounting holes in the Pi 5 PCB, the active cooler connects to the fan power connector. This fan cools extremely well and stays relatively quiet. We tested the original Pi 4 case fan and it sounded like tinnitus. The active cooler is quiet and only ramps up

the speed when it really needs it.

At the time of writing, the official case and the active cooler are the only known cooling products for the Raspberry Pi 5. Reusing Raspberry Pi 4 coolers and cases will be problematic, largely down to the placement and removal of ports on the Pi 5. The composite jack has been removed and we now have two



Connecting the Raspberry Pi 5 to a USB SSD adaptor is simple and can be done for less than £20.

camera/screen connectors in its place. The USB and Ethernet ports have been swapped to their original locations (the Pi 4 changed them around) and this makes nearly all Pi 4 cases incompatible. The swap to micro HDMI ports means that Pi 3 cases are also incompatible. We did have a go at retrofitting a 52Pi Ice Tower cooler to the Pi 5. It worked, but we wouldn't suggest that you try doing this.

Long story short, buying the active cooler is a must, wait a while for a case.

Operating system

At the time of writing the Raspberry Pi 5 has just been launched, and for those who were quick off the blocks, their pre-orders are in the post. You need an operating system for your new Pi 5 and right now there are two choices. The official Raspberry Pi OS Bookworm (see *review, page 42*) release is a great desktop OS. Now based on Wayland, it feels slick and very grown-up, but it does have a couple of issues. Firstly, a change to the underlying Debian 12 distro means that installing Python modules via *Pip* is not possible at an OS level. PEP 688 details that we can no longer do this and we now need to use Python Virtual Environments (*venv*) in order to install our favourite Python modules. If you want to install Python modules at an OS level, use the *Apt* package manager. The second issue lies in the GPIO. Yes we can use the GPIO with Python, but we

YOU NEED

- > **Pi 5**
- > **USB 3 to SATA/NVMe adaptor**
- > **NVMe or SATA SSD**
- > **Pi 5 Active Cooler**



Using SD Card Copier, we can clone our Raspberry Pi OS Bookworm install from microSD to a much faster USB 3/SSD combo.

have to use GPIO Zero or get comfortable with `I2C`. `RPi.GPIO`, the venerable Python module that powers the majority of HATs and Pi projects, no longer works with the Raspberry Pi OS. If you need to use a HAT, use an older Pi with Bullseye for now. Software support for your favourite HATs is coming, but still some time off.

The second OS choice for the Raspberry Pi 5 is Ubuntu 23.10, which has always been a good second-place OS in our eyes. It has always looked good and performed decently, but it had its own GPIO issues. With Ubuntu 23.10, we still have some HAT compatibility issues, but the GPIO is supported. Add the power of the Raspberry Pi 5's quad-core 2.4GHz CPU and Ubuntu 23.10 runs slickly.

Emulation-based OSes are still a work in progress; Recalbox has created a test version for the Pi 5, but your favourite emulator may not be supported yet.

Overclocking a Pi 5

We've overclocked every model of Raspberry Pi and that now includes the Raspberry Pi 5. We've taken the CPU to 3GHz and the GPU (VideoCore VII 800MHz stock) to 1.1GHz!

Your results may vary, because the odds of the silicon lottery are never evenly stacked.

To overclock, firstly you need active cooling, so buy the active cooler before you attempt this.

In a terminal, open `/boot/config.txt` using Nano:

```
$ sudo nano /boot/config.txt
```

Scroll to the end of the file and add a new comment to identify where the overclock lines are:

```
# These lines overclock the Pi 5
```

Now add these two lines to overclock your Raspberry Pi 5 CPU and GPU to 2.8GHz and 900MHz respectively. Why are we not going straight to 3GHz and 1.1GHz? We need to creep up to those speeds and ensure that our system is stable before we attempt a faster overclock.

```
arm_freq=2800
```

```
gpu_freq=900
```

Save the file by pressing Ctrl+X, then Y and Enter. Then reboot the Pi for the change to take effect.

To test the stability of the overclock, install `Stress`:

```
$ sudo apt update && sudo apt install stress
```

Run a five-minute stress test on all the CPU cores:

```
$ stress --cpu 4 -t 300
```

Open another terminal and use the `vcgencmd` command to monitor the CPU speed every second.

```
$ watch -n 1 vcgencmd measure_clock arm
```


Open a further terminal and use `vcgencmd` to monitor the GPU speed every second.

```
$ watch -n 1 vcgencmd measure_clock core
```

If you would like to learn more about `vcgencmd` `00`, you can head over to www.tomshardware.com/how-to/raspberry-pi-benchmark-vcgencmd, where you'll find detailed information.

Bootling from USB

The Raspberry Pi 5's microSD card speed is phenomenal, but perhaps you would like capacity and speed at a lower price point? USB-to-SATA/NVMe adaptors provide easy access to cheap storage. We



The *Raspi-config* tool has more options than the Raspberry Pi Configuration GUI tool. It can help us tweak the system to meet our needs.

used an Argon One NVMe adaptor and a spare NVMe drive. Booting from microSD, we opened a terminal and launched the *Raspi-config* tool:

```
$ sudo raspi-config
```

Go to Advanced Options > Bootloader Options and select the latest Bootloader. This flashes the latest bootloader to the Raspberry Pi 5's firmware. You then need to reboot.

Go back into *Raspi-config* and select Advanced Options > Boot Order and select USB Boot. Don't reboot but close *Raspi-config*.

Connect your USB 3 drive to a USB 3 port and go to the main menu. Select Accessories > SD Card Copier.

From the app, select the source drive – this is your microSD card. The destination is the USB 3 drive. Click Start and the microSD card install is copied to the USB drive. This may take some time. When done, power off the Raspberry Pi 5, remove the microSD card and then power up again.

Your Raspberry Pi 5 just got a little bit faster and cooler, and has access to plentiful storage, and all it took was a handful of lines of code, some button presses and a few extra pounds. **LXF**

» POWER SUPPLIES

Any electronic device is useless without power and the Raspberry Pi 5 needs a good power supply. The official Raspberry Pi 5 PSU is the 27W USB-C power supply, and this custom-designed PSU has all the power we could need. Seriously, this USB-C PD PSU provides 5.1V at 5A – that's more power than a Steam Deck PSU (5V at 3A). The Raspberry Pi doesn't use anywhere near this much power, but when we start to add accessories, such as the hotly anticipated M.2 HAT, we will start to draw more power from the supply.

We recommend the official 27W PSU for the Raspberry Pi 5, though you can use the older Raspberry Pi 4 official PSU, but that is limited to 3A and you may start to see the dreaded 'lightning bolt' on your OS. Other power supplies can be used, but we would stick with supplies that deliver a true 5V output. USB-C PD supplies can be tricky. It is normally the responsibility of the device to negotiate the correct voltage, but if there is a mistake, your precious Raspberry Pi 5 may feel the sting, and so will your wallet.

You can also provide power via the GPIO, but we suggest that only experienced makers attempt this. You need to provide the correct voltage via a regulator and ensure there is enough current.

» **GET YOUR Pi FILLING HERE** Subscribe now at <http://bit.ly/LinuxFormat>

Saving Private Files using your own NAS

Christian Cawley needs to back up his PC. So of course, he chooses to do it with a Raspberry Pi.



OUR EXPERT

Christian Cawley is that bloke who buys every single Raspberry Pi model, does something really cool with it, then puts it in a drawer until the next one comes out.

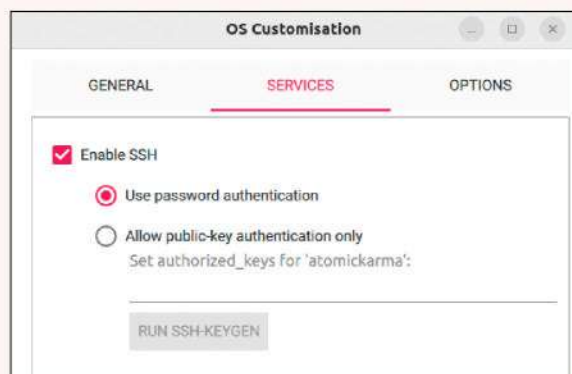
Making and keeping your own backups is more secure and reliable than banking on a third party and an internet connection. While any PC can be configured as a backup machine, a more compact solution like the Raspberry Pi keeps such a set up low-powered and almost infinitely flexible.

Different backup types are available, with a full backup, specific file/device backups, and incremental (those that simply make a copy of any files that have changed) all possible. Using a Raspberry Pi as a cloud-accessible NAS is straightforward, but choosing the right software for the task is less obvious.

What is a NAS?

Network-attached storage is a file server that can be accessed across a network over Wi-Fi or Ethernet. Unlike a *Samba* server, which relies only on file browser access, a NAS supports browser and client app access.

Various well-known cloud platforms bring NAS-like functionality to your PC, but have a key shortcoming: cloud platforms aren't hosted by you. Your data is held remotely, by a third party that you may or may not pay monthly. If any data is considered unwelcome, it can't be retained in the cloud. Furthermore, remote server issues can impact access to your own data. So, it makes sense to manage your own data. With the right



Whichever Raspberry Pi OS you use, save time by enabling SSH at the installation stage in Raspberry Pi Imager.

software, reliable network and fast, plentiful storage, a Raspberry Pi 3 or later can make an excellent NAS.

Dropping the Box

Over the years, this author has moved from one cloud storage solution to another, using them to sync key files on his computer. Back when he was a Windows user (forgive him), he retained data on SkyDrive, which later became OneDrive. Spotting some performance issues with the Microsoft 'solution',

» UNDERSTANDING DOCKER CONTAINERS

Docker support for Raspberry Pi has improved considerably over the past couple of years, with straightforward, easy installation and management that is almost on a par with desktop systems.

If you're unfamiliar with *Docker*, it's worth spending a few moments considering how this software works, and the impact of using it on a Raspberry Pi with external storage attached.

Docker is a virtualisation system that enables you to run one or more applications in an isolated environment. These applications are prepackaged and ready to run, and are known as containers. They feature all the prerequisites and supporting software for the application, so little or no additional installation is required. Containers are also lightweight and can be easily

shared or removed. They can even be moved to a new location, with no impact on performance and stability.

What all of this means is that when it comes to downloading your *Docker* container(s) on a Raspberry Pi, they can be on any mounted drive. This is really useful because the YML file that is part of the *Seafile* configuration apparently doesn't support specifying a

path to an external drive. It also doesn't support symlinks. Therefore, the solution is to install the *Seafile* container on the external drive.

In terms of performance, the hit is negligible. You're switching from using the SD card to using a USB drive, preferably an SSD or mechanical HDD, rather than USB flash. So, you get faster data transfer and the SD card enjoys a longer lifespan.

he then moved to Dropbox, then on to Box, before switching back to Dropbox via Google Drive.

While these are all relatively reliable (the quantity of files stored on them varies your mileage), they all have one thing in common: your data is beyond your control. For redundancy, this might seem like a good idea – after all, if your property goes up in a ball of flames, your data is almost certainly gone, too. But retaining control of your data is a fundamental right.

Your data might include banking and ID records. It could be creative works – everything from reports, projects and novels-in-progress to family photos and videos, music projects, game saves, recipes and even downloaded audiobooks and ebooks.

There is also the issue of bloat. Popular cloud storage provides browser-based document viewing and editing. This is unnecessary in many cases, especially if all you need is a flexible backup solution.

Ten years ago, bringing this data under your control relied on a dedicated NAS or an old PC reconfigured as such. Now, thanks to the power of later Pi models and affordable storage, a NAS can be set up within an hour.

Which NAS solution?

There are several options available with which to build a NAS. *Nextcloud* (<https://bit.ly/LXF310-nextcloud>) is a popular cloud storage service that offers a downloadable home server setup. While widely supported, there is a lot that this service offers that is simply irrelevant to this project (such as a daily dashboard, cloud documents and more). *Nextcloud* has its uses, but it's overkill for hosting backups.

Preparing this guide, we also spent a lot of time with *OwnCloud* (<https://bit.ly/LXF310-owncloud>). Probably too much. Early on, it became apparent that this solution was going to get messy (certainly as a manual installation). PHP incompatibilities (the project is locked to PHP7, with no sign of a change in policy), Apache issues, and problems with *MySQL/MariaDB* make this an option to avoid. While fully compiled images are available, *OwnCloud* has far more features than are necessary for a flexible NAS backup project.


Seafile (<https://bit.ly/LXF310-seafile>) offers a far more reliable setup experience. It is also suited to use as a file server. While the alternatives attempt to offer everything Dropbox, Google Drive or even your employer's intranet/SharePoint server provides, *Seafile* plays to its key strengths and purpose.

Fitting out your Pi NAS

Strictly speaking, this process requires a compact, single-board computer and a hard disk drive, preferably an SSD with a SATA-to-USB interface. As long as the required OS and software runs on the board, you can get it up, running and backing up your PC.

For the purposes of this tutorial, we have set up a Raspberry Pi 4 (8GB) with a 1TB SSD. This is mounted in a Desk Pi Pro case, and relies on a USB 3.0 interface with the SATA slot. A single power supply to the case handles both the Pi and the storage device.

Your own requirements may differ, or you may not have a Raspberry Pi 4. This shouldn't be a problem, as



any Pi 3 or later should reliably handle incremental backups fine. You could potentially build a NAS backup solution using a Pi Zero and a cranky old 5.25-inch mechanical disk drive, but we'd advise against this.

Attempting this project with a Raspberry Pi 5 and its PCIe interface delivers superior read/write speeds, of course, regardless of whether you use an NVMe or SATA device. At the time of writing, however, this isn't an option – not as an off-the-shelf purchase, at least.

Prep your Pi

For the best results, begin with a fresh Raspberry Pi OS installation. This is a headless project, so you can use a lite OS. If you're using the *Raspberry Pi Imager*, select Choose OS > Raspberry Pi OS (Other) for lite versions.

Next, connect an external USB disk drive, preferably SSD or HDD. This needs to be mounted and formatted (if not already done) with the ext4 filesystem. Chances are your drive is already formatted as FAT or NTFS; by default, ntfs-3g should be installed on Raspberry Pi OS, making it simple to mount the drive. If not, install with:

```
$ sudo apt install ntfs-3g
```

Next, reconnect (or power cycle) the disk and check its name with `lsblk`. The typical result of this is `/dev/sda` but ensure this is the case before proceeding – you don't want to format the wrong device!

Next, run `Fdisk` to list partitions on the device:

```
$ sudo fdisk -l /dev/sda
```

(Note that the `sdx` name is for SATA devices. An NVMe device will be named `/dev/nvmeNnN`, where N is a number, 1-9.)

Use `Fdisk` to format the drive:

```
$ sudo fdisk /dev/sda
```

Hit `m` to check the commands – you can use this reference at any point. To format the disk, use `D` to delete any partitions (ensuring first you're not losing any important data), then select `G` for a new empty GPT partition table.

Next, hit `N` to create a new partition, then `1` to use the entire disk as a single partition. You need to enter the partition size; for example, a 1TB disk is around 930GB, so use `+930G`.

When you're done, tap `P` to display the partition table, then `W` to write the changes to the disk.

Next, format the partition as ext4:

```
$ sudo mkfs.ext4 -F /dev/sda1
```

With the disk partition formatted, create a mount point for the device:

```
$ sudo mkdir -p /mnt/seafile
```

Mount the partition:

```
$ sudo mount /dev/sda1 /mnt/seafile
```

The final thing to do with the disk is edit the `fstab` file, ensuring the device mounts when the Pi reboots.

Start by checking the UUID of the partition:

```
$ blkid /dev/sda1
```

The Docker software downloads Seafile in multiple parts, with progress bars demonstrating completion of each pull.

QUICK TIP

Setting this project up over SSH saves time, so remember to enable the interface – and configure the username and password when installing your Raspberry Pi OS.

QUICK TIP

You can access your Raspberry Pi file server from any local device, including your smartphone browser. Consequently, you can start backups from your phone while you eat dinner, watch TV, or read a book in bed.



» GET TO SAVE EVERY LXF ISSUE! Subscribe now at <http://bit.ly/LinuxFormat>

```

atomickarm
GNU nano 7.2
db:
image: mariadb
container_name: seafile-mysql
environment:
  - MYSQL_ROOT_PASSWORD=password
  - MYSQL_LOG_CONSOLE=true
volumes:
  - ./data/mariadb:/var/lib/mysql
networks:
  - seafile-net

memcached:
image: memcached
container_name: seafile-memcached
entrypoint: memcached -m 256
networks:
  - seafile-net

seafile:
image: seafileltd/seafile-mc
container_name: seafile
ports:
  - "8080:80"
volumes:
  - ./data/ann:/shared

```

The YML file provides opportunity to specify database passwords and establish a username and password for the Seafile admin account.

QUICK TIP

While you can connect a standard external hard drive via USB, an all-in-one case with space for a Raspberry Pi and an SSD is more flexible and power-efficient. A USB drive is also an option, but you're likely to find that read and write speeds are slower for larger capacity devices.

Make a note of the UUID, then open **fstab**:

```
$ sudo nano /etc/fstab
```

The UUID should be entered at the end of the file on a new line: **UUID=[YOUR UUID] /mnt/ext4 defaults,auto,users,rw,nofail,noatime 0 0**

Here, **ext4** refers to the filesystem, **defaults** specifies kernel default options: **auto**, **users**, **rw**, **nofail** and **noatime**. That final flag improves system performance. The two 0s at the end are largely unused and refer to system dumps – a zero value ensures this is not enabled.

Set up your server

Before installing *Docker* and the *Seafile* container, you need to install the *Nginx* web server software. This is straightforward:

```
$ sudo apt install nginx
```

The default instructions for this do not require HTTPS, because you want to get things up and running first. However, this can be set up on *Nginx* with *Python* and *Certbot*:

```
$ sudo apt install certbot python3-certbot-nginx
```

This is optional, however, and can be done later if required, once you've established whether *Seafile* is the solution you need.

Seafile requires an *Nginx* profile. This is provided in a CONF file in *Nano*:

```
$ sudo nano /etc/nginx/sites-available/seafile.conf
```

Enter this text:

```

server {
    server_name localhost;
    location / {
        proxy_pass http://localhost:8080;
        proxy_set_header Host $host;
        proxy_set_header X-Real-IP $remote_addr;
    }
}

```

Where you see **localhost** in the code, change this to the IP address of your Raspberry Pi (if you're using *Certbot* and a URL, you change it to the URL).

The next step is to copy the file to the *Nginx* **/sites-enabled/** directory with a symlink, then restart *Nginx*:

```
$ sudo ln -s /etc/nginx/sites-available/seafile.conf /etc/nginx/sites-enabled/seafile.conf
```

```
$ sudo systemctl restart nginx
```

(If implementing *Certbot*, enter **sudo certbot** at this stage, following the instructions.)

Next, it's time to install *Docker*. You need both the main *Docker* software and the *Docker Compose* tool, which handles the YML file. This makes managing any containers easier.

```
$ sudo apt install docker.io docker-compose
```

Wait for installation, then run **hello world**:

```
$ sudo docker run --rm hello-world
```

You should see the 'Hello from Docker!' message at this point.

Install and configure Seafile

As mentioned, installing the *Seafile* container on the external drive ensures that device is used for backups. Start by creating a **seafile-server** directory and changing directory into it:

```
$ sudo mkdir /mnt/seafile/seafile-server && cd /mnt/seafile/seafile-server
```

Next, create a **docker-compose.yml** file:

```
$ nano docker-compose.yml
```

You can find the contents of this file on the *Seafile* website (<https://bit.ly/LXF310-seafileyml>). Ensure you change **MYSQL_ROOT_PASSWORD** and **DB_ROOT_PASSWORD** to the same value. You also need to set the **TIME_ZONE** (ETC/UTC if in the UK) and set a **SEAFILE_ADMIN_EMAIL** and **SEAFILE_ADMIN_PASSWORD**. (If using HTTPS, set **SEAFILE_SERVER_LETSENCRYPT** to **True** and specify a **SEAFILE_SERVER_HOSTNAME**. Otherwise, leave these fields as per the defaults.)

Press **Ctrl+X** to exit, **Y** to confirm the save, then **Enter**. Now download the *Seafile* container with:

```
$ sudo docker-compose up -d
```

Wait for this to complete. You'll find it is relatively quick, even over a slower internet connection. Once done, input the **logs** command to ensure the container is running:

```
$ sudo docker logs seafile
```

Your *Seafile* file sync is up and running. Check it on your local network by opening a browser and entering the Pi's IP address.

You're now ready to make a backup and sync it to your PC (see *walkthrough*, opposite). **LXF**

» HOW BIG SHOULD YOUR STORAGE BE?

Although later Raspberry Pi models can handle a 1TB microSD card, wear and tear on such a device when storing a regularly updated backup reduces its lifespan considerably. Given the importance of the data you're archiving, it makes sense to add reliable storage to your Raspberry Pi.

But what you're planning to back up has a direct impact

on the storage you connect to the Pi. Full PC backups take up more space than, say, a regular archival of the images you have stored on your PC. Both might be considered equally important, but they typically have different storage requirements.

Then you need to consider speed. As mentioned (see *Understanding Docker Containers*, page 46), an SSD

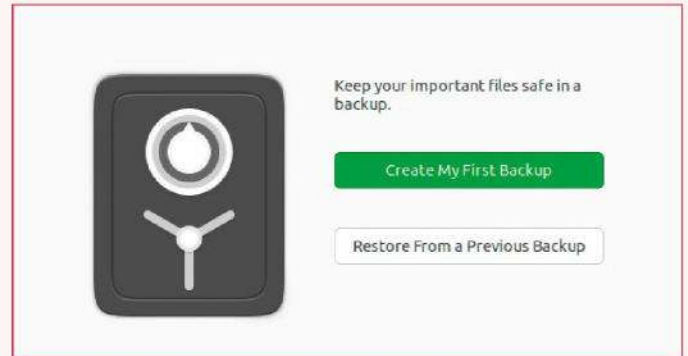
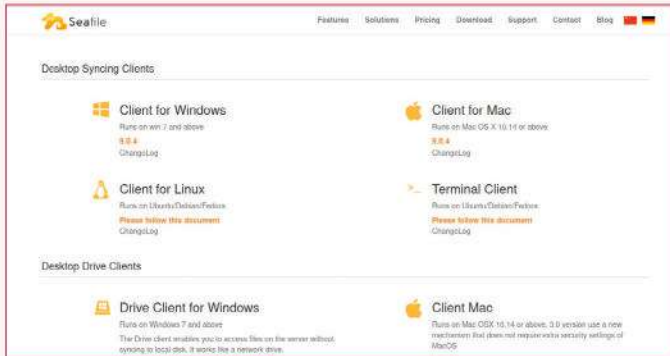
is more efficient, facilitating faster backups. Additionally, with its lack of moving parts, an SSD uses less power than an HDD. However, you might have a spare HDD (as we did) sitting around doing little more than collecting dust.

When it comes to price, an SSD is typically slightly more expensive than an HDD. Low-capacity drives are priced similarly, but as

you climb the ladder of gigabytes and terabytes, the prices deviate noticeably.

At the very least, your destination drive should offer enough capacity for a single backup archive. Keep to your budget, but appreciate there is going to be a trade-off: an expensive, efficient SSD could prove cheaper in the long term than a slower, power-demanding HDD.

SYNC YOUR PC BACKUPS TO SEAFILE

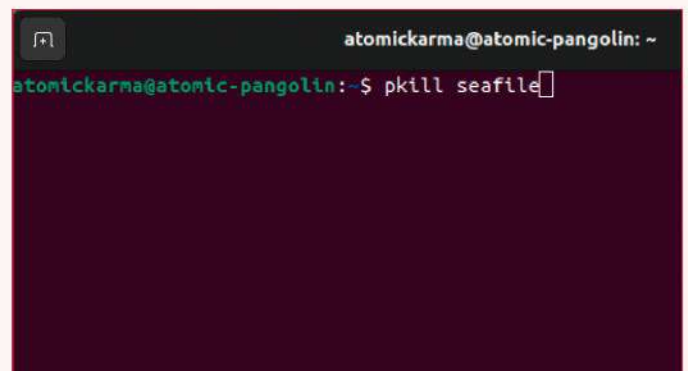


1 Set up the Seafile client

To use *Seafile* as a file server, you need to install the client that matches your desktop operating system. Apps for Linux, Windows and Mac desktops can be all downloaded from www.seafile.com/en/download. Connect to the client over your local network using your Raspberry Pi's IP address and the credentials you established in the `docker-compose.yml` file.

2 Create a backup

Various data types can be archived in a backup. This example uses a full disk backup generated with the *Gnome Backups* tool on Ubuntu, but any backup file can be synchronised with *Seafile*. For simplicity, use your operating system's built-in backup tool (third-party solutions add overheads when it comes to recovery).

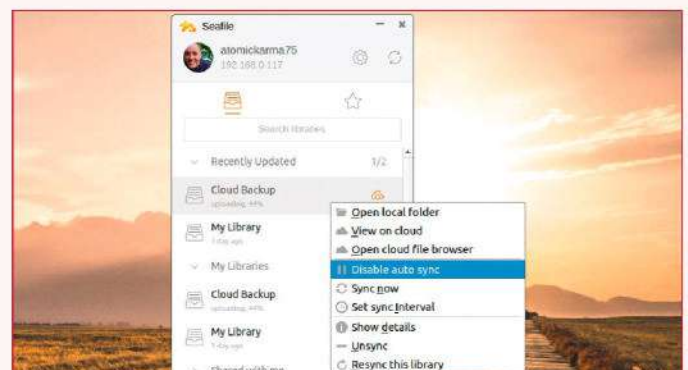
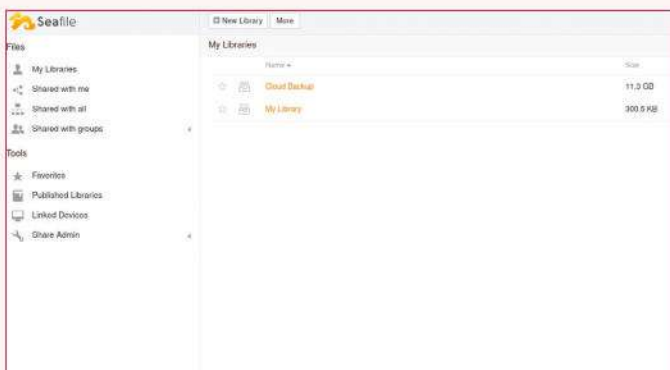


3 Save the backup to your sync folder

Connect the client to *Seafile* using the Pi's IP address, and the username and password credentials you established in the `docker-compose.yml` file. When the client connects to the server, you can send the backup. In the *Seafile* client, click Select, browse to the backup directory, and click Open.

4 No sync? Restart

Network issues can disrupt a *Seafile* sync. If it appears to have stopped or failed, with no upload or download rate displayed in the bottom-right corner of the client, restart it. Clicking X to close keeps it running in the background, so use `pkill seafile` in your terminal to restart.




5 Review the synchronised data

Whether you are synchronising for the first time or have done so many times, it is always worth reviewing synchronised data. You can view all data uploaded to *Seafile* in your web browser. Just input the Raspberry Pi's URL, sign in, and click My Libraries to review the data.

6 Disable reverse synchronising

Because synchronising can ruin a backed up file (say, if a failed backup process damages the local copy), it is wise to disable automatic syncing. In the *Seafile* client, right-click the synced directory and click Disable Auto Sync. This prevents bad backups overwriting good ones.



Tails Tails

Nate Drake explores how to configure the latest version of Tails for maximum security and convenience.

Total online anonymity is impossible, but if used correctly, the Tor network, or dark net, comes pretty close. This is the thinking behind Tails (The Amnesic Incognito Live System), a Debian-based security-focused Linux distro.

Since 2009, the good people of the Tails Project have hitched their star to Tor's wagon by making sure that whenever you use the OS, all internet connections are routed through the dark net.

If you're new to the dark net and want to familiarise yourself with how Tor is configured to conceal your identity, it's now time to stop. Head on over to www.torproject.org and return here when you're done.

If, on the other hand, you're happy the dark net can provide the online safety you need, you can do far worse than use Tails. Instead of offering a standalone browser, you have an entire OS geared towards keeping you safe online. Connections outside Tor are blocked (with one exception, as you'll learn) and the system can be run in amnesiac mode. In other words, it can load into your device's RAM, then erase critical data on shutdown. This makes it almost impossible to recover your data using current digital forensics.

Tails also supports a persistent mode, whereby you can create an encrypted partition on a USB stick to

store vital settings, like browser bookmarks and Wi-Fi passwords. In this guide, you'll learn how to set up and configure this, as well as the pros and cons of doing so.

What about the Tor Browser?

First, let's address the elephant in the room. The Tor Project does indeed offer the *Tor Browser*, which by default doesn't store your browsing data and also routes all connections through the dark net.

Tails itself uses an implementation of the browser (currently v12.5.6) with the uBlock extension, so the Tor Project must be doing something right.

Still, Tails offers a number of advantages over just installing the *Tor Browser* on your regular machine. Firstly, even browsers run in private mode like the *Tor Browser* can leave traces; such programs save files temporarily to your machine's hard drive, so they can be recovered to build up data on your browsing habits.

The *Tor Browser* also doesn't use any encryption, so anyone with access to your user account can open it to view your bookmarks and/or view your connection settings. The same applies to any files you download.

When it comes to sharing files, if you've prepared them on your personal device, there's always a risk they could contain identifying metadata, even if sent

via the dark net. Tails makes it easy to remove this information, plus files created in live mode naturally don't contain any metadata from your hard disk.

Running Tails in live mode via USB or DVD also makes it very unlikely that any malware installed on your device's hard drive can compromise your safety, given that the OS is loaded into your RAM.

Finally and most importantly, if you use the *Tor Browser* at the same time as other internet software that isn't connected to the dark net, a sophisticated adversary can use this to build a picture of your browsing habits.

Getting started

If you've decided that Tails can offer the privacy you need, open your regular web browser and head on to <https://tails.net/install/linux/index.en.html#install>.

You also need a blank USB stick of at least 8GB. Follow the steps to download the Tails USB image (currently 1.3GB), then install it via your chosen disk utility. If you're using a Debian-based distro like Ubuntu, *Gnome Disks* is perfect for this. If it doesn't come bundled with your OS of choice, just open the terminal and run `sudo apt-get install gnome-disk-utility`.

From here, it's just a matter of a few mouse clicks to 'restore' the disk image to the USB stick. Remember that all data already on the USB drive will be erased.

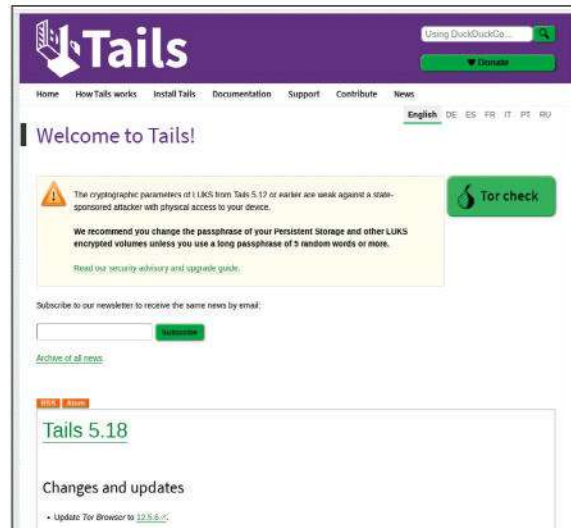
Depending on your device, you may need to use the Boot key or configure your UEFI to make sure Tails boots. Save any open files and restart your machine.

Assuming that Tails loads correctly, you'll see a welcome screen asking you to choose your language, keyboard layout and number formats. Click Start Tails at the top-right to load the desktop.

The default desktop looks very spartan but will be familiar to Debian users. (Incidentally, the most recent version of Tails is based on Debian 12 Bookworm.)

Start by clicking on Network Settings at the top-right and connecting to your chosen Wi-Fi network. By default, the OS is in amnesiac mode, so doesn't store settings like this, nor any files you download.

Once you're connected to a Wi-Fi network, hit Activities at the top-left and then choose Tor



Tails comes with the Tor Browser and defaults to the developer's page. Click Tor Check to ensure the connection is set up correctly.

Connection. This opens a Tails-specific wizard to begin configuring your connection.

Building bridges

If you don't mind anyone knowing you're using Tor, you can choose Connect To Tor Automatically. However, if you're living in a jurisdiction where Tor connections are monitored and/or blocked, you need to use a bridge. This makes it far more difficult for someone monitoring your internet traffic to know you're using Tor.

Having derided the standalone *Tor Browser*, the most recent version has a number of features that the *Tails Browser* doesn't. In particular, it makes it much easier to connect to the Tor network via a bridge, as it comes with a number of preconfigured ones. It also supports the Snowflake pluggable transport, which operates similarly to a bridge but more efficiently.

If you choose Hide To My Local Network That I'm Connecting To Tor, you need to manually configure an obfs4 Tor bridge. This is a nuisance, as you need the details for a valid bridge and then have to enter it into the Tails connection wizard via a QR code or type it out.

In our tests, we decided instead to use *Tails Unsafe Browser*, which connects directly to the internet to visit

» A-PEELING ANONYMITY WITH ONIONSHARE

OnionShare is an open source tool that lets you securely and anonymously share files as well as chat with friends using the Tor network. Tails comes with this program preinstalled – it can be used in both live and persistent mode.

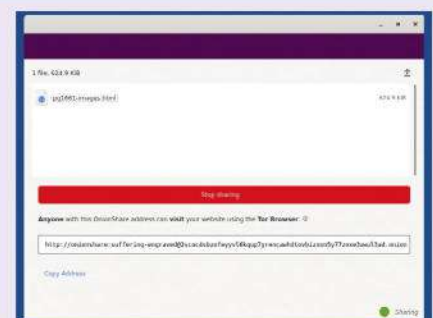
To get started, just click Activities and search for *OnionShare* via the search bar. The interface couldn't be simpler. To anonymously send a file that is already on your system, just click Share Files. Next, drag and drop your file or files to the main window, then choose Start Sharing.

OnionShare now generates a link that anyone can use to download your file(s) using the *Tor Browser*. (This is

necessary because the addresses use the .onion TLD.)

If you have a website you want to publish, *OnionShare* supports setting up your own Tor hidden service. Download your site to a suitable folder on Tails, start *OnionShare* and choose Publish Website. Click Start Sharing to generate a unique .onion address where others can access the site.

You can also use *OnionShare* to receive files from others directly to your computer. To do this, start the app, then choose Receive Files. *OnionShare* generates a unique .onion address to enable you to receive files to your computer, but only while it's



running. Ensure you read the warnings about downloading and opening untrusted files before using this feature. If you've enabled persistence, make sure that any files received are stored in your **Persistent** folder before rebooting.



<https://bridges.torproject.org>, in order to obtain a bridge address to copy and paste into the Bridge field. You can then click Connect To Tor. If you experience any errors at this stage, check that the system clock is showing the correct time.

If your local network has a captive portal, you may also need to use the *Unsafe Browser* to go online before the Tor connection can be established.

Assuming all goes well, you can now click Start Tor Browser and close the wizard.

Take some time to familiarise yourself with the *Tor Browser*. If a website is available as a Tor hidden service (using the .onion TLD), the browser asks whether you want to switch. We recommend using .onion addresses where possible, because this keeps your traffic entirely inside the dark net, making it harder to trace and intercept.

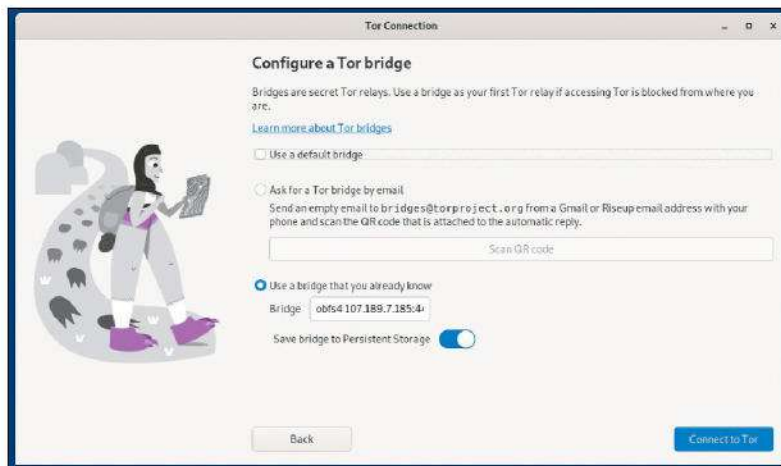
Persistence pays off

There's a lot to be said for continuing to use Tails in its current configuration. By default, your account has no admin rights, making it much harder for malware to make a purchase in your system.

All trace of your browsing activity is lost on reboot and Tails also has a utility to delete personal data from RAM upon shutdown, making it more resistant to cold boot attacks than other operating systems.

Still, this isn't very practical for day-to-day use. Each time you restart, you need to enter the Wi-Fi password and configure Tor bridges, and the browser doesn't save any bookmarks. Even if you use some of the handy bundled apps such as *LibreOffice* to create documents, they're lost after you shut down. This also applies to any extra programs you install.

If you're booting Tails from a USB stick, however, you can enable persistent storage. This enables you to store certain files and settings between reboots on a partition encrypted using LUKS. Crucially, if you set this



With persistent storage enabled, you can preserve Tor bridge settings across reboots. Visit <https://bridges.torproject.org> for obfs4 bridge addresses.

up, you don't have to unlock the encrypted partition each time, so can still use Tails in amnesiac mode.

Before going any further, we recommend a thorough read of <https://tails.net/contribute/design/persistence/> to check this is the right option for you.

To get started, simply click Activities, then Persistent Storage. Click Continue to enter your password. The wizard suggests a high-strength passphrase, though we recommend using *Diceware* (<https://bit.ly/diceware>) to generate your own.

Next, the persistent storage wizard asks you to specify which features you want to keep across reboots. By default, the **Persistent** folder is enabled, which you can access from the Places menu at the top-left. You can also enable persistence for network settings, such as Wi-Fi passwords, Tor bridges (if you have one), browser bookmarks and additional software just by clicking the rocker switch next to each category.

If you do enable persistence, apart from making a backup, you need to be clear about what is and isn't saved. For instance, while your browser bookmarks exist across reboots, you still need to configure the browser to show the bookmarks toolbar each time.

Adding to the long tail

Tails comes with a good crop of preinstalled apps, such as the *LibreOffice* suite and *Mozilla Thunderbird*, but if you've enabled persistence, you may want to add a few more from Debian repositories.

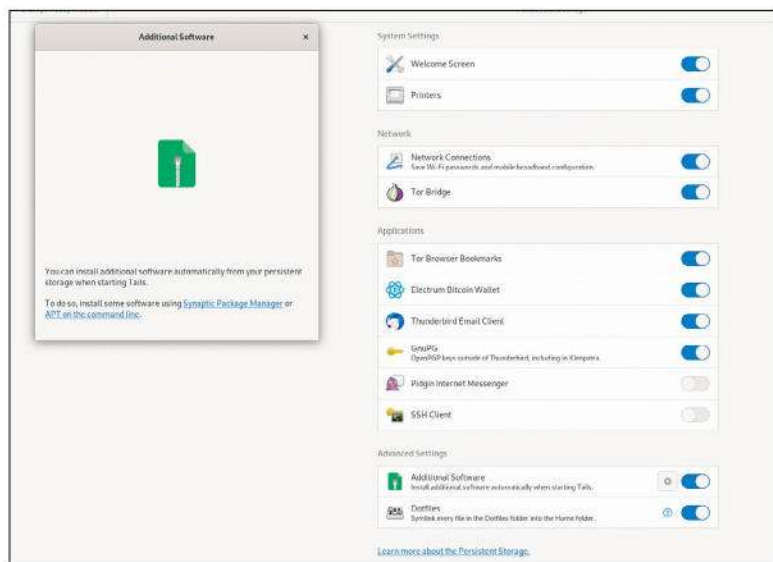
This is straightforward after you've enabled the additional software feature. You can either launch the *Synaptic* package manager to browse available apps or install via **apt-get** using the command line.

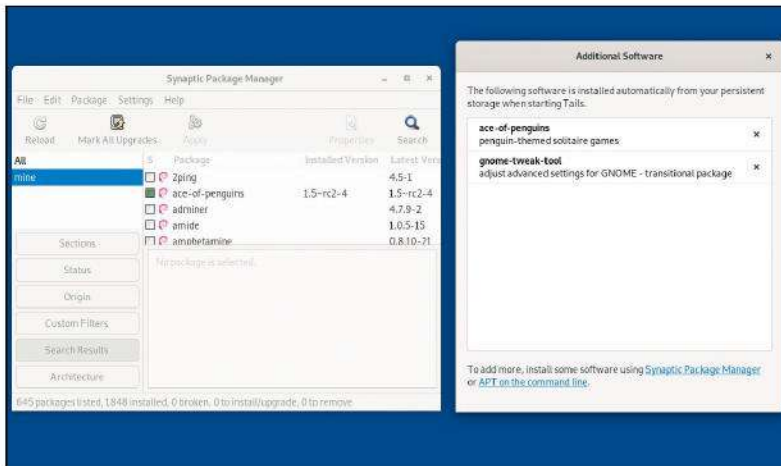
The only caveat is that, by default, Tails doesn't give you admin rights. To fix this, be sure to restart Tails and then, once on the welcome screen, click on the + under Additional Settings. From here, you can double-click Administrator Password and set one to allow you to perform admin tasks.

If you choose to install additional software, Tails gives you the choice to do so for this session only or to add it permanently. If you choose the latter, it's reinstalled from your persistent volume each time you reboot, so allow time for this.

As for the *Tor Browser*, reinstalling software does not necessarily mean all settings are carried across

Tails is amnesiac by default but can be configured to preserve certain settings and files in an encrypted volume.





You can use Synaptic or the terminal to add new programs. Tails supports installing programs once only or across reboots.

reboots. For instance, after installing *Gnome Tweaks*, we found that each time Tails restarted, we had to go back into the utility to enable window centring.

Tails of woe

The Tails Project website has a helpful warning page explaining that the OS is “safe but not magic” (<https://tails.net/doc/about/warnings/index.en.html>).

For instance, while using Tails in live mode protects against most forms of malware, it doesn't offer much protection against more sophisticated attacks, such as compromised firmware or a hardware keylogger. Needless to say, you should only run Tails on devices you trust and control.

There aren't currently any known viruses that can infect a Tails installation, but to be safe, make sure you fully power devices before inserting the USB.

If you choose to enable persistent storage, it'll be obvious to anyone with access to your USB drive that it contains encrypted data, because once plugged into any Linux device, it prompts immediately for the passphrase. Provided the password is strong enough, they won't be able to access your private information. But if your device is seized or lost, you won't be able to either, which is why we recommend making a backup (see box, right).

The OS also can't protect your personal information if you use it to share files with metadata that can be traced to you. On the plus side, Tails does include *MetaData Cleaner* to scrub such data from files before you send them. Set up separate email or chat accounts to share files or, better yet, use *OnionShare*, which comes preinstalled in Tails (see box, page 51).

In this guide, we've assumed you're using a wireless connection, which can itself be sniffed and exploited. You can eliminate this risk by connecting to your router or modem via Ethernet cable, either directly or with a home plug device.

If your local network or ISP has active policies to block Tor connections, you should first try connecting via an unlisted obfs4 bridge. Some users have also had joy combining Tor connections with a VPN or SSH connection (see <https://gitlab.torproject.org/legacy/trac/-/wikis/doc/TorPlusVPN>).

However, this would require extra work for Tails, both in terms of setup and preserving your VPN/SSH

configuration and credentials between reboots, so we only suggest this for experienced Linux users.

Turning Tails

Tails can provide an excellent way to stay safe online. You can make life more convenient for yourself (if a little less private) by enabling persistence to save data to an encrypted partition, provided you're willing to do some fine-tuning after each reboot.

We've already discussed the OS's advantages over the regular *Tor Browser* but there are other ways to access the dark net.

Other well-established secure operating systems include Qubes and Whonix, though these can be more difficult than Tails to set up and configure correctly.

You could use a live bootable USB of most versions of Linux and junction all network connections via Tor, but this is less anonymous, because Tails' uniform design makes it far harder to distinguish individual users via methods such as browser fingerprinting.

Head over to <https://tails.net/doc> and read through the online documentation to decide whether this is the right privacy solution for you. **LXF**

» PERSISTENT PROBLEMS

While enabling persistent storage allows you to preserve some configurations and personal files across reboots, there's the risk that something can happen to your USB stick.

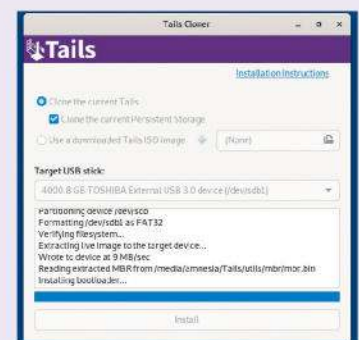
Fortunately, Tails makes it easy to run backups of this data. Just insert a different external drive into your device, choose Activities and click the icon for *Tails Cloner*.

The cloner offers the choice to copy your entire Tails installation or just the persistent storage. If you choose to back up the persistent storage, the utility prompts you to choose a passphrase for the encrypted partition on the new device.

Choose your external drive from the Target USB Stick drop-down, then click Install. Read the dire warning that all data on the target drive will be lost, then click Delete All Data And Install.

Tails Cloner then partitions/formats the drive and copies the partition(s) to the new drive.

In future, if you want to update your backup, simply plug in the



backup drive, then go to Applications > System Tools > Back Up Persistent Storage.

You can also do this manually via the command line. Load Tails with persistence enabled and insert your backup drive. Click on the encrypted volume and enter the passphrase to unlock. Open the root terminal and run:

```
rsync -PaSHXv --del /live/
persistence/TailsData_
unlocked/ /media/amnesia/
TailsData
```

Once the command completes, it displays a summary of all copied data. Learn more at https://tails.net/doc/persistent_storage/backup/.

Safeguard your secrets

Shashank Sharma gets a taste for 007's life by reading fiction and occasionally sending encrypted messages to his unsuspecting brother.



**OUR
EXPERT**

Shashank Sharma is a trial lawyer in Delhi, India. He's been writing about open source software for 20 years, and lawyering for 10.

Some messages are just too private to be shared directly with the recipient. Although the need might not arise often, it's imperative that you know the right encryption tool if you want to share messages privately. But the word encryption conjures a jumble of other words, such as algorithms, ciphers, salts, public and private keys, and so on, which might seem like a little too much trouble to master when all you want to do is organise a surprise party for a beloved boss's retirement.

Enc bills itself as a modern and friendly alternative to *GnuPG*, which has long been the torch bearer for encryption across Linux and Unix systems. Released under the MIT licence, *Enc* is not as robust as *GnuPG*, and does not aim to be a stand-in replacement for it in every instance. In fact, *Enc* is designed with the classic Unix philosophy of 'doing one thing well' in mind. To that end, *Enc* can be used to encrypt messages. And that's it. Of course, the message could be stored in a text file and encrypted using a myriad of ways, such as passwords or keys.

While *GnuPG* is part of the software repositories on most distributions and is often part of the default installation, you have to install *Enc* manually. Thankfully,

» GETTING HELP

The *Enc* project is limited in its functionality and doesn't aim to compete with *GnuPG* in terms of supported features. This enables the project to provide documentation on all of its limited feature set. The GitHub page walks you through the different aspects of using *Enc*, without inundating you with tech jargon.

Unlike most other GitHub-hosted projects, which have an easily accessible Issues page, *Enc* doesn't list this for some reason. You have to manually point your browser to <https://github.com/life4/enc/issues> if you wish to report an issue.

In any event, you can access help for each of the subcommands from the terminal itself, by using `enc <command> --help`. For instance, if you need to see an example usage of the `enc armor` command, or view the available flags, you can run `enc armor -h`.

Almost all the subcommands have a shorter alias, usually their first letter, but you can verify this by accessing the specific help page. For instance, to generate a private key, you would use `enc key generate > private.key`, but you can also use `enc k g > private.key` instead.

```
linux@linuxformat.com:~$ enc key generate --help
Generate new private key

Usage:
  enc key generate [flags]

Aliases:
  generate, create, g

Examples:
  enc key generate > private.key

Flags:
  -b, --bits int           size of RSA key in bits (default 4096)
  -c, --comment string     a note to add to the key
  -e, --email string       your email address
  -h, --help               help for generate
  -n, --name string        your full name
  -t, --ttl string         validity period of the key. Can be a date (2020-12-30)
                           or duration (4y30d, 24h) (default "1yr")
  -t, --type string        type of the key (default "rsa")
linux@linuxformat.com:~$
```

Almost all commands in *Enc* have a shorter alias to help you save time by not having to type the complete command.

the process is quite straightforward. If you already have Go installed on your system, you can run the `go install github.com/life4/enc@latest` command to install the latest release.

Alternatively, you can download the latest source tarball from the project's GitHub page (<https://github.com/life4/enc/releases>), run the `tar -xvf enc_<version>_<release>.tar.gz` command, and move the extracted binary to any directory in your `$PATH`:

```
$ cd Download/projects
$ mkdir enc
$ mv enc_1.1.2_linux_x86_64.tar.gz enc
$ cd enc
$ tar xvf enc_1.1.2_linux_x86_64.tar.gz
$ ls
enc enc_1.1.2_linux_x86_64.tar.gz LICENSE README.md
$ chmod +x enc
$ sudo mv enc /usr/local/bin
```

Unlike most other utilities, which extract files into a dedicated directory, you have to manually create a directory and move the tarball to it, before extracting files. Next, we make the `enc` binary executable, and finally we move into a directory in the `$PATH`, in this case the `/usr/local/bin` directory.

Tell me a secret

The basic workflow when working with *Enc* is that you use the `enc encrypt` command to encrypt your message or text file with a password. When you want

to read the file or message, you have to decrypt the encrypted binary file using the same password and the `enc decrypt` command.

Let's create a text file with a message that we wish to keep secret from prying eyes with the `touch secretfile.txt` command. You can now use your favourite text editor to add some content to this newly created file, and save it.

```
$ cat secretfile.txt | enc encrypt --password 'some-clever-password' > encryptedfile.bin
```

In this example, we've used the `enc encrypt` command to safeguard the message using a password. To decrypt the file, we now use the `enc decrypt` command, as you'd expect:

```
$ cat encryptedfile.bin | enc decrypt --password 'some-clever-password'
```

```
This is a test text file.
```

```
It will be deleted once the tutorial is done.
```

As you can see, the content of the encrypted file is displayed on screen when you provide the correct password. The *Enc* utility reports an error if you type the wrong password:

```
$ cat encryptedfile.bin | enc decrypt --password 'same-clever-password'
```

```
Error: cannot decrypt message: gopenpgp: error in reading password protected message: wrong password or malformed message
```

Because of the way the content of a file or message is directed to `enc encrypt` and `enc decrypt`, the *Enc* utility can only work with text files or text messages.

You don't always have to place your secret message in a text file before redirecting it to `enc encrypt`:

```
$ echo 'this message will not self destruct no matter how long you wait' | enc encrypt --password 'genius-password' > encryptedfile.bin
```

```
$ cat encryptedfile.bin | enc decrypt --password
```

```
'genius-password'
```

```
this message will not self destruct no matter how long you wait
```

By default, *Enc* displays the decrypted message on the screen, but you can divert it to an output file instead if you prefer with the `cat encryptedfile.bin | enc decrypt --password 'genius-password' > decryptedmsg.txt` command.

Some apps and utilities don't support sharing encrypted files. If you ever have to pass an encrypted file through such an application, you can disguise the encrypted binary as a text file, using `enc armour`:

```
$ cat encryptedfile.bin | enc armour > armoured.txt
```

```
$ cat armoured.txt
```

```
-----BEGIN PGP MESSAGE-----
```

```
Version: enc 0.1.0
```

```
Comment: https://github.com/life4/enc
```

```
wy4ECQMikWbWtibgOXbgl5zqljHp500RXIKPC5HZ4x12
OnEBPndO49+iotHu+B4GozpFjBfBETgMEQMfT9ymz7F8
VPXy3JhziE9QbHH+ToyLQspUNfMQ3MMYkKSBbVilj44
OxpMhOH5x0SRojSD1U/stnKf0g==
=X2Ly
```

```
-----END PGP MESSAGE-----
```

Note that the subcommand is `armour`. You have to decrypt the armoured file to read the secret message:

```
$ cat armoured.txt | enc decrypt --password 'genius-password'
```

```
this message will not self destruct no matter how long you wait
```

The project's GitHub page (<https://github.com/life4/enc>) has some useful pointers on how you can get past having to type your secret passwords as part of the terminal command, which can then be read by accessing the Bash history.

Working with keys

The other aspect of encryption is that it can be used to allow others to send you secret messages that can only be decoded by you. This is done by sharing with the world a public key, such as on your website or on your social media page. All secret messages intended for you are encrypted using this public key. Such messages can only be read by decrypting them using your private key.

To generate a private key, you must run `enc key generate > private.key`. This private key is valid for one year by default, and comprises a host of information such as your name, email address and so on.

You can provide custom values to the key, such as email, duration and so on. Run `enc key generate --help` to view all the definable parameters for the key.

The public key associated with the private key so generated can be extracted with the `cat private.key | enc key public > public.key` command.

To encrypt a message with the public key, you must run the `cat file.txt | enc encrypt --key public.key > encrypted-with-key.bin` command. This encrypted file can be read with the `cat encrypted-with-key.bin | enc decrypt --key private.key` command.

For additional security, you can also lock your private key with a password. This is recommended, in fact, otherwise anyone with access to your system and private key can read all private messages.

Run `cat private.key | enc key lock --password 'password-for-private-key' > passwordlocked.key` to password-protect your private key. To decrypt a message using the locked private key, run `cat encrypted-with-key.bin | enc decrypt --key passwordlocked.key --password 'password-for-private-key'`.

Refer to the Getting Help box (*opposite*) for instructions on how you can learn more about *Enc* and its various subcommands. **LXF**

```
linuxlala@playground:~$ enc --help
A user-friendly CLI tool to work with PGP keys:
create, add, list, encrypt, decrypt, sign, verify signatures.

Usage:
enc [command]

Available Commands:
armor      Convert the message (or signature) from binary to text
completion Generate the autocompletion script for the specified shell
dearmor    Convert the message (or signature) from text to binary
decrypt     Decrypt the message
encrypt     Encrypt the message
help       Help about any command
key        Operations with key
keys       Operations with key ring
remote     Operations with remote key servers
sig        Operations with signatures
version    Print the version number

Flags:
-h, --help  help for enc
```

QUICK TIP

If you're unfamiliar with shell redirection using `|` (pipe) and `>` (redirect output), a quick read-up would help. You can look it up online, or run the `man shell` command, and read through the Redirection section.

Credit: <https://github.com/tesseract-ocr/tesseract>

Dump your paper docs with perfect OCR

Nick Peers reveals how to extract editable text from images and printed materials with the help of optical character recognition software.

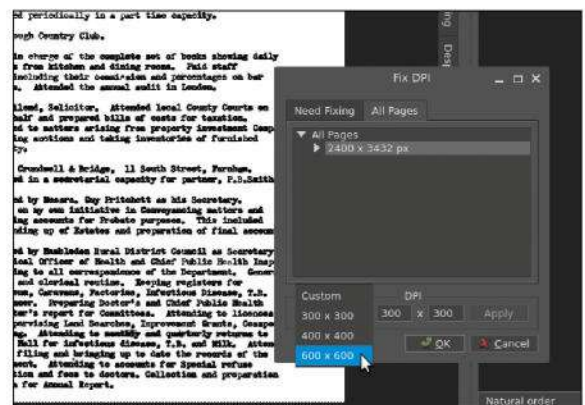


**OUR
EXPERT**

Nick Peers loves his family history, and loves it even more now he can convert family biographies into editable text files without manually transcribing them first.

Anyone who's been faced with the arduous task of transcribing text from either printed material or a digital scan of printed text will no doubt have heard of optical character recognition (OCR) technology. OCR is one of the earliest examples of machine learning, whereby a computer model is trained to recognise shapes on a digital image and translate those shapes into text characters. Once the shape of each letter is identified and translated into editable text, words followed by sentences, paragraphs and entire tracts of text can be extracted from the digital scan.

OCR has roots going back to the 1980s, and while commercial engines perform increasingly miraculous conversions – not just on typed text, but also handwriting – open source engines continue to develop alongside them. Linux is blessed with several OCR engines, all with roots in commercial products, but now open sourced and completely free to use. The best known of these – which we'll focus on in this tutorial – is Tesseract (<https://github.com/tesseract-ocr/tesseract>), a command-line OCR engine that can be used on its own or paired with a number of graphical front-ends to perform OCR across a variety of usage scenarios, from extracting editable text directly from scanned documents to converting everything from



If your scanned image isn't perfect, use ScanTailor Advanced to make it clearer for Tesseract to read and translate.

ApplImage and Snap – but you can also install it via its own repository to ensure it's the latest version (5, as opposed to 4 in most universal repositories):

```
$ sudo add-apt-repository ppa:alex-p/tesseract-ocr5
$ sudo apt update
$ sudo apt install tesseract-ocr
```

This installs both the Tesseract engine, plus trained data enabling it to recognise English text. You can add more languages using `sudo apt install tesseract-ocr-lan`, substituting *lan* with the relevant country code, such as *fra*, *spa* or *deu*.

With the command-line engine installed, you can now perform OCR from the command line on a wide range of files – prior to 2016, Tesseract only worked with TIFF files, but now it can handle most popular formats, including PNG, GIF and JPEG. Basic command-line usage is as follows:

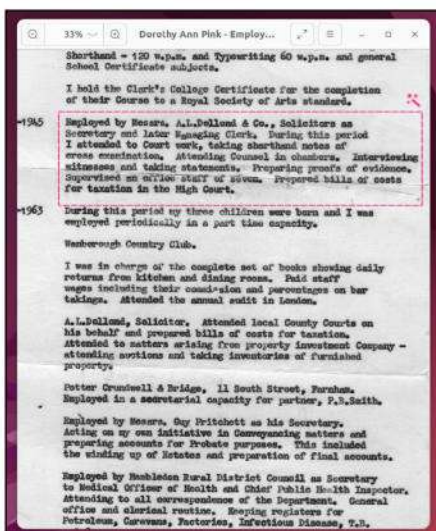
```
$ tesseract image.jpg textfile
```

Full documentation can be found under Command Line Usage at the program's online documentation (<https://tesseract-ocr.github.io/tessdoc/>), but instead of persevering with the terminal, let's examine a selection of GUI tools that provide a user-friendly front-end to Tesseract while also expanding its capabilities.

One front-end to rule them all

If you're looking for a more user-friendly way to perform OCR on a wide range of filetypes as well as

Choose which part of the screen to capture and NormCap OCRs its contents and places them on the clipboard.



PDFs and image files to screen grabs and image-based subtitle tracks in media files, too.

Before going further, check the box opposite for a quick look at Tesseract and two of its main open-source rivals – note, you can install all three at once and try different ones to see which produces the best results.

Marks, set, scan!

The obvious place to start is by installing the underlying Tesseract OCR engine. It exists in various forms – including standalone

printed documents (via your scanner), there are two major apps worth considering. On the surface, *GImageReader* (<https://github.com/manisandro/GImageReader>) seems the more obvious choice given it's actively maintained and capable of working with a wide range of input sources, including screenshots and the clipboard's contents.

The big issue with *GImageReader* on older Ubuntu 22.04-based systems (including Mint) is that it doesn't support Tesseract 5.x. In the absence of a Snap or recent Flatpak build, you need to remove both it and the PPA from your system, then install the older Tesseract 4.x engine from the system repos. This results in less accurate OCR, but if that's not a deal-breaker – or you're running Ubuntu 23.04 or 23.10 – check out the annotation (right) for a quick rundown of how it works.

Thankfully, there's an alternative for older systems that does work with Tesseract 5: *OCRFeeder*. Unlike *GImageReader*, the app is almost dormant, with just the occasional update; nevertheless, it works fine with Tesseract 5.x, and can easily be installed through your software store or repo (`sudo apt install ocrfeeder`).

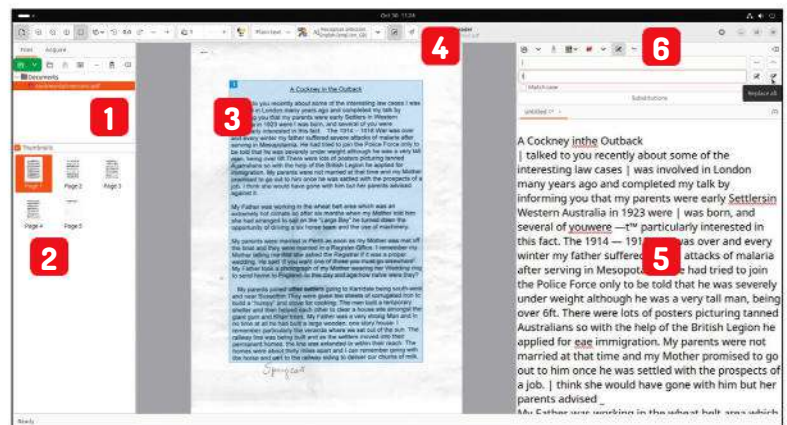
Once in place, open *OCRFeeder* from its shortcut on the launcher or Start menu (look under Graphics in Mint). You're presented with a two-pane screen, ready for you to import your first file or scan your first page. The step-by-step guide (*over the page*) takes you through the basic process of importing images or PDF files, and converting them into an editable document. When exporting, you have a choice of four formats: ODT (*LibreOffice Writer*), PDF and HTML (which attempt to preserve the original layout), and Text to Simple, which creates a plain text document, albeit with plenty of blank space.

Review and edit

After your file has been generated, it pays to open it in its native application and examine the results – these depend on the quality of the original image or scan as well as the font used. It's likely that the transcription isn't perfect, so take the time to read through it and make any corrections.

It's worth noting that you can also make corrections to your scanned text in *OCRFeeder* itself – use the Text tab under Text Properties to review and edit if required. If Tesseract does a particularly poor job, consider

OCR WITH GIMAGEREADER



1 Input file(s)

Use this pane to import any files you wish to OCR, or switch to the Acquire tab to grab from your scanner.

2 Thumbnails

A list of all the separate pages within the current document are shown here – click one to jump to it.

3 Preview window

Click and drag your mouse around the portion of the image that you wish to convert to OCR.

4 Perform OCR

Click here to run Tesseract on the selected text, or click the down arrow to access more options.

5 Edit text

Once OCR is performed, the recognised text is pasted here. Review and make corrections if required.

6 More options

You'll find buttons for saving your file as text, plus performing other formatting, such as stripping out line breaks.

trying a different OCR engine on particularly badly handled parts of the image – after installing another engine (see the box, below), select it from the drop-down menu and click the OCR button to see if produces better results.

Next to the Text tab, you'll see two further tabs: Style enables you to choose font, size, alignment and spacing for the currently selected portion of text, while Misc lets you confirm which language is being used to perform the OCR itself.

On some scanned pages, you may find flecks of dirt are picked up as separate text boxes (bounded in blue). To exclude these from your document, select each one in turn and click Type > Image – they should turn green.

Improving OCR quality

A major reason why Tesseract doesn't produce perfect results is the quality of the original scan or digital

QUICK TIP

When it comes to converting handwritten notes to text, you have to trust your notes to the cloud – Google Drive is one option whereby you can import scanned images or PDFs of your notes and have them converted automatically to editable text.

» TRY DIFFERENT ENGINES

We've focused on the Tesseract OCR engine in this tutorial, but it's not the only option. Two alternatives are also available and can be installed alongside Tesseract for comparison purposes. In fact, several front-end tools like *OCRFeeder* are designed to work with multiple OCR engines – they detect any you've installed to enable you to switch between them

depending on which one you want to use.

The first is *Cuneiform*, a mix of freeware and open source code, so is designated non-free. You can install version 1.1.0 independently through your universal repository – `sudo apt install cuneiform` – as it's not been updated in a very long time, other than to address minor compatibility issues. The

command-line tool does have options (`--dotmatrix` and `-fax`) for improving accuracy on faint or fragmented text.

The other is *GOCR* (<https://jocr.sourceforge.net>). It is also no longer under active development; its last major update was in 2018. Again, install it through your repos (`sudo apt install gocr`).

In most cases, however, Tesseract is likely to be the

only engine you need. Make sure you have the latest build (5.x) installed where possible, as this is the best developed and most accurate one available. If you come across a front-end that doesn't work with the latest version, see if you can source a Flatpak or Snap version, which should bundle a compatible (if less accurate) version of Tesseract for you to use.

QUICK TIP

To open **NormCap** via a shortcut, go to **Settings > Keyboards > View and Customise Shortcuts in Ubuntu**. Select **Custom Shortcuts** and click **Add Shortcut**. See <https://dynobo.github.io/normcap/#faqs-start-via-cli> for the commands; **Ctrl+Shift+PrtScr** is usually available as a shortcut.

image. If you're working with a less-than-perfect image (clear, black text on a paper-white background), take the time to try to improve the quality of your scan or digital image, either using your scanner's own controls or through an image editor such as **GIMP**.

When scanning, boost the resolution as high as you can if results are disappointing – try 600dpi rather than 150 or even 300dpi. Also make sure the pages are as square as possible to prevent the text warping. If you're unable to prevent this – or have obtained the images from another source – try installing a post-processing tool called **ScanTailor**, which can resolve many of these problems.

There are several builds available, but **ScanTailor Advanced** (<https://github.com/4lex4/scantailor-advanced>) is one of the best, and can be installed through Flatpak or Snap:

```
$ sudo snap install scantailor-advanced
```

Once installed, launch the program and import your scanned images. Work through the options on the left to fix such problems as skewed pages while boosting contrast, increasing DPI and converting to black and white to bring out the text more clearly. Once done, files are exported in TIF format, which can then be

imported into any of our featured OCR tools with hopefully better results.

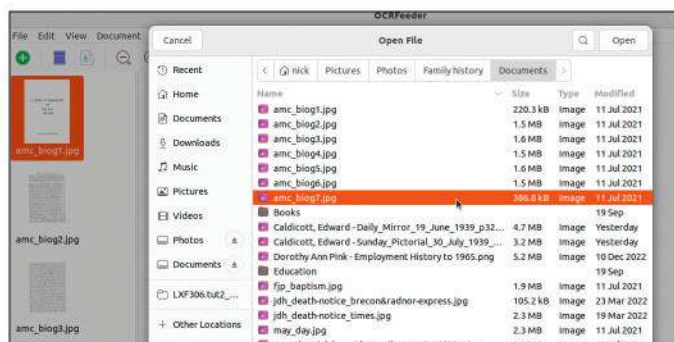
Capture screen text

What happens when you come across some text on screen that you can't select – for example, text within embedded images in web pages or error messages from non-selectable dialog boxes? The answer lies with a screenshot tool with built-in OCR capabilities. There are several tools available, but one common issue they all have is working well with the newer Wayland desktop manager.

If you're willing to use the X11/Xorg display manager instead (you can log off your current session, then click your username followed by the settings button to switch), you're almost spoiled for choice. One of the best is **NormCap** (<https://dynobo.github.io/normcap/>), which can be run as a portable AppImage or installed through Flatpak:

```
$ flatpak install flathub com.github.dynobo.normcap
```

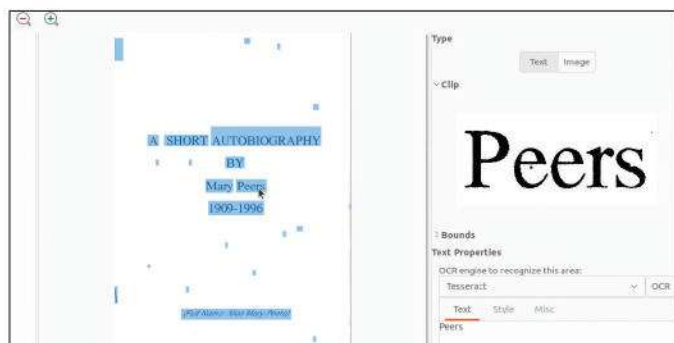
Once installed, launch it via its shortcut (or from the command line with `flatpak run com.github.dynobo.normcap`) and wait for a pink settings button to appear on-screen beneath the menu bar.

CONVERT IMAGES TO TEXT QUICKLY**1 Import existing files**

Open the File menu and choose **Add Image** or **Add Folder** to add all supported image files (TIFF, PNG, JPEG or GIF) within the selected folder. Repeat for any more images you wish to add or choose **Import PDF** to import a single or multi-page PDF.

**2 Rearrange order**

Multi-page PDFs are converted to individual JPEG images and imported accordingly. Once all your pages are in place, use the left-hand pane to drag and drop files to change the pagination order, or select an unwanted page and choose **Edit > Delete**.

**3 Detect text**

Once your pages are lined up and in the correct order, choose **Tools > Recognise Document** to scan each page in turn for recognisable text. The status bar keeps you apprised of its progress. Once complete, areas recognised as text are highlighted in blue.

**4 Review and export**

Click on a blue section and you should see the detected text appear in the right-hand pane. You can make manual tweaks to the text from here, change the style, and even rescan the selected area with another OCR engine. Once happy, choose **File > Export** to export your text.

A crosshair will also appear – simply click and drag this around a portion of the screen containing text and *NormCap* grabs the screen, OCRs any detected text using your underlying *Tesseract* installation, and places it on the clipboard ready for pasting elsewhere.

Click the settings button to reveal more options – including a choice of parse and raw capture modes. Parse mode – the default – attempts to detect (and preserve) sentences that break over several lines as well as paragraphs and URLs. It can even attempt to detect multiple email addresses and turn them into a comma-separated list, suitable for using in CSV files. Raw mode simply returns the text as it's detected – including any line breaks.

You'll see that only English is supported as a language by default, but by choosing Add/Remove... from the app's menu, you can quickly download more, then switch between them.

By default, *NormCap* quits once you've captured your text and transferred it to the clipboard; choose Keep In System Tray to add a shortcut to the menu bar that remains visible at all times, or check out the Quick Tip (left) to find out how to launch it via a universal keyboard shortcut whenever it's needed.

NormCap does work with Wayland, but only in Gnome 43 or later, which rules out Ubuntu 22.04 for now. Instead, it's easier to switch to Xorg whenever you anticipate having to capture text from a screenshot.

Image to text

If you've ever attempted to convert a file to MP4 format in *Handbrake* while including any subtitle files, you may have occasionally come across a situation where the subtitles can only be burned directly on to the image itself. This is down to the fact that certain streams use PGS subtitles, which render subtitles as images rather than text.

These subtitles need to be converted to a text format, which of course involves OCR. Several tools exist that enable you to do this, but the obvious choice (KDE's *Subtitle Composer*, <https://subtitlecomposer.kde.org>) requires you to manually teach the editor by laboriously going through the feed telling it what each character is.

The easiest way to run this is via the portable Applmage provided at <https://subtitlecomposer.kde.org/download.html> – after launching, click the Open Subtitle button to select the file containing your subtitles. You need to select All Image Subtitles from the Files Of Type drop-down menu to reveal your file (typically, the original MKV or VOB file). Select this, then click Open.

Make sure the correct stream/language has been selected (typically there's only one). You'll see a list of post-processing options are all ticked by default – for now, leave them as they are and click OK. Next, resize the import window to make things easier to read.

You'll see the first subtitle line appear in the box with the first character highlighted in red. Type this into the box and click OK or press Enter – it moves on to the next character. Repeat for each unique character (A-Z, a-z, 1-0, symbols) – as you progress, *Subtitle*

» ADD EDITABLE TEXT LAYER TO PDFS

One major shortcoming with *Tesseract* is that it doesn't natively support PDF files. Thankfully, if you're looking for a way to OCR PDF documents from the command line, *OCRmyPDF* is the obvious tool. It works in tandem with *Tesseract*, converting PDFs into image files (one per page) that are then run through *Tesseract*'s OCR engine before being stitched back together into a single PDF. This new PDF contains two layers – the existing image layer, plus an additional text layer that is both editable and searchable in PDF viewers and editors.

Installation is straightforward either through your repo (`sudo apt install ocrmypdf`) or Snap (`sudo snap install ocrmypdf`), the latter guaranteeing you have the latest version. Once installed, navigate to the directory containing your target PDF file, and issue the following: `$ ocrmypdf filename.pdf output.pdf`

This takes your input file (**filename.pdf**) and generates a PDF/A file called **output.pdf** (obviously, you can rename this to whatever you like). PDF/A files are designed for long-term storage, so embed any required fonts for display purposes and cannot be encrypted.

There are numerous extra flags available (`--output-type pdf` for standard PDF files, for example), and you can review these in detail at the program's online documentation (<https://ocrmypdf.readthedocs.io>) in the Cookbook section.

```
nick@Ubuntu: ~$ cd Documents/
nick@Ubuntu: ~/Documents$ ocrmypdf cockney-dah-version.pdf newpdf.pdf
Scanning contents: 100%|██████████| 5/5 [00:00<00:00, 260.54page/s]
Start processing 4 pages concurrently
OCR: 100%|██████████| 5.0/5.0 [00:26<00:00, 5.23s/page]
Postprocessing...
PDF/A conversion: 100%|██████████| 5/5 [00:00<00:00, 5.81page/s]
Recompressing JPEGs: 0image [00:00, ?image/s]
Deflating JPEGs: 100%|██████████| 5/5 [00:00<00:00, 33.10image/s]
JBIG2: 0item [00:00, ?item/s]
Optimize ratio: 1.03 savings: 2.5%
Output file is a PDF/A-2B (as expected)
nick@Ubuntu: ~/Documents$
```

If you don't need a GUI tool, but want to perform OCR on a PDF, *OCRmyPDF* extends *Tesseract*'s capabilities to do just that.

Composer starts to skip most (but not all) previously recognised characters.

If you make a mistake, just click Previous Symbol and make your correction before clicking OK. If you come across certain types of punctuation like a colon or semi-colon, or the letters i or j, only one part of the symbol is selected – click + next to Symbols to increase the number to 2 and both should be selected, enabling you to train the system to recognise that symbol; if they're not, click Next Symbol to skip to the next one. If two characters are mistakenly recognised as one, simply type both symbols before hitting Enter.

Look out for stylised text (italics, bold) – use the formatting controls above the OK button to instruct *Subtitle Composer* to recognise this as styled text, and it ensures the text-based subtitles honour that styling.

If this is too much hassle, check out the Quick Tip (right) for an alternative that uses *Tesseract* to speed up the process, but even here you need to go through the subtitles once complete as mistakes do occur. **LXF**

QUICK TIP

If you're happy to install Mono and work with a tool that was written specifically for Windows, visit www.nikse.dk/subtitleedit/help#linux for a guide to installing *Subtitle Edit*, which should process image-based subtitles much more quickly.

» **STICK WITH PAPER-BASED LXF!** Subscribe now at <http://bit.ly/LinuxFormat>

WordPress security, events and users

Michael Reed concludes his overview of setting up and maintaining a WordPress site with some finishing touches and a look at plugins.


**OUR
EXPERT**

Michael Reed has been fiddling with websites since the dial-up days and the 10MB of free hosting space that his ISP gave him. He doesn't like to bring up such hardships now. Much.

After installing WordPress and customising it in parts one and two of this series, we're ready to start extending its functionality with the plugin system. By the end of this series, you should be ready to handle every major aspect of managing a WordPress site. We're now going to add the finishing touches so the site is ready to go on the internet for use by the general public.

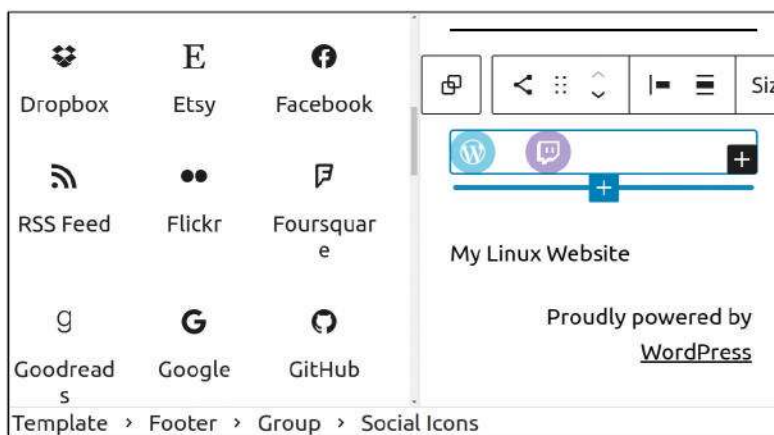
We'll assume that you have a working WordPress installation that you can access via a web browser, and the WordPress installation can be located on professional hosting or installed on your own server. Last month, we customised the appearance of the site, but you don't need to have followed any of those procedures to try out the ideas that we're delving into this month.

Managing plugins

Plugins are the principal mechanism for extending the functionality of your WordPress website, and there are nearly 60,000 to choose from. To begin the process of choosing and adding plugins, make sure you are signed into your WordPress website as an administrator (go to [http://\[URL of website\]/wp-admin](http://[URL of website]/wp-admin)). If you're signed in but not in the Dashboard, use the link at the top-left of the window to navigate to it.

We'll be spending a lot of time this month on the Plugins administration page, accessible from the side menu of the Dashboard. This page lists the installed plugins and also enables you to activate, deactivate and delete plugins. Deactivation is sometimes required when troubleshooting, because it enables you to eliminate a given plugin when you're trying to track down the cause of a problem on your site. Many of the more complicated plugins also have a settings option that takes you to the settings page for that particular plugin. Other plugins add the link to their settings page to the sidebar of the Dashboard.

This page also shows which plugins have an update available; do this by clicking on Update. However, you



WordPress has an impressive collection of social media profile icons ready for linking.

can enable the auto-updates feature on your plugins, and you can do this per individual plugin by clicking on the Enable Auto-Updates link, but you can also carry out bulk management operations. To do this, click on the empty checkbox before the plugin name for multiple selection or click on the empty checkbox at the top or bottom of the entire list. There's a number of actions that you can carry out that are selectable from the Bulk Actions drop-down menu, and it can be a real time saver when managing a lot of plugins.

Freemium isn't free

WordPress plugins are either free, premium (paid for) or freemium (free version of a commercial plugin). Freemium plugins typically have a subset of the full features, and they are often not open source. This might sound off-putting, but a freemium plugin is often a good choice within the WordPress plugin ecosystem.

By reading the description and reviews, you might find that you can live with the restrictions of a freemium plugin; often, premium features are for corporate users and less important to typical users. If you do need the more advanced and extensive features, it might be worth considering investing in the full version of a premium plugin.

The main reason that the free version of a freemium plugin is often such a good option is that it means that the development is funded by sales of the full version,

QUICK TIP

When installing plugins, don't install unnecessary ones, as they can slow WordPress down and introduce the possibility of conflicts.

tested with your version of WordPress, plugin hasn't been updated for a while have been abandoned. It's also an indication the plugin isn't regularly updated, which



Every Post and Page has a full set of social media sharing icons at the bottom, thanks to the AddToAny plugin.

and this means that resources can be directed at maintaining and updating a plugin over the long term.

Installing a plugin

Installing a plugin follows much the same process as adding a theme, which we covered last month. On the Plugins page, click on Add New to begin the process. This takes you to a page listing available plugins – there is a set of tabs at the top of the page to select between categories such as Recommended and Popular. Frankly, we don't find these categories to be particularly useful due to the massive number of available plugins with overlapping functionality.

As often as not, you're better off following a plugin recommendation from a trusted source or searching for a plugin by keyword within the WordPress interface. The search box for doing this is on the top-right of the Plugins page, and although it is labelled Keywords, you can search for plugins by name as well.

There are a few things to look out for when choosing a plugin from the search results. Start by checking out the rating and reviews. Obviously, avoid any with a very low rating. The panel for an individual plugin in the search results also shows whether or not the plugin has been tested with your version of WordPress. There is a chance that a plugin might work if it hasn't, but there is some risk to site stability and security. It's also an indication that the plugin isn't regularly updated, which is a must in the long term.

Add social networking icons

There are two types of social media icons you might want to add to your WordPress content: icons that link to your own social media accounts (such as Facebook or Instagram) and those that facilitate the sharing of your content by readers of your site. In the latter case, they are posting the links on their own social media pages. This draws new readers to your WordPress site.

We can add the first kind of social media profile links without resorting to plugins. For example, let's add some social media links to the footer of the site. Go to Dashboard > Appearance > Editor to enter the site editor. You might remember from last month that the representation of the front page of the site that takes over the right area of the editor is not actually editable itself; so, click anywhere on it to begin editing.

Making sure that the Template tab in the right-hand sidebar is selected, click on Footer under Areas. Click on the footer area of the site, above existing content, such as the site name, as this is where we want the social media icons to be placed. This causes the

floating toolbar to appear. Click the three dots and select Add Before. This makes a + icon appear; click on it and select Social Icons (you may need search for it).

This causes another + icon to appear, and clicking on this causes another searchable box with some icons to appear. With this, you can start adding the social media icons. Click on one, such as the Facebook icon, and this is added to the footer of your website. Click on the icon and add the URL of your Facebook page. Repeat this for all of your social media sites. Phew, that was a lot of clicking on + icons, but it's fairly intuitive once you're used to working with that area of WordPress. Finally, don't forget to click on Save (in the right-hand sidebar) to make the changes permanent.

Add sharing icons

WordPress has built-in facilities to add links to your social media pages, but if you want to give readers the ability to share your content in a convenient way, you have to use a plugin. There's quite a lot of choices but we settled on a plugin called AddToAny.

We searched for AddToAny in the WordPress interface (Dashboard > Plugins > Add New). It was shown to be compatible with our version of WordPress, and it had more than 400,000 installations and a solid 5/5 rating from more than a thousand votes. A click on the Details link confirmed that it was what we were after, and the user reviews themselves made it clear

QUICK TIP

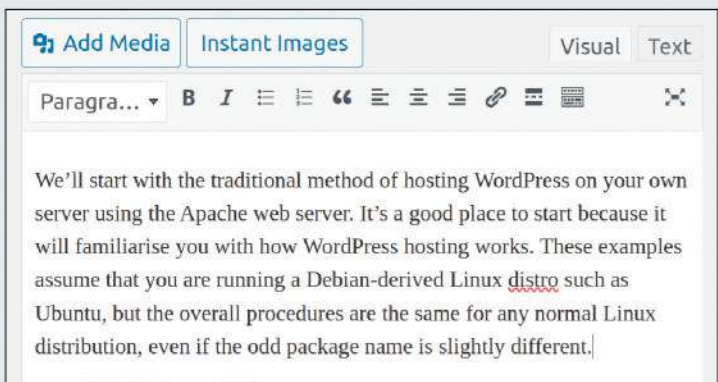
Be careful when searching for plugins on the web as some blog authors are actually connected to the plugin developers, and this isn't always made apparent.

» BACK TO THE CLASSICS

WordPress's built-in, block-based editor, Gutenberg is modern and efficient, but what if you want to work with your WordPress site without learning new skills? The classic editor that you might already be familiar with is still maintained and updated.

The first plugin to add is Classic Editor. As it's an official plugin supported by WordPress, you shouldn't expect to have any problems when running it on an otherwise modern WordPress setup. It gives you back the old Post and Pages editor with a toolbar at the top of the window. There are, however, some limitations when it comes to editing the site layout using the classic editor. As administrator, you can set the classic editor to be the default for all users or you can allow each user to choose, but once you've started using either editor, that editor subsequently has to be used on that post or page.

Classic Widgets is another plugin that restores the classic functionality to modern WordPress because it enables you to place widgets, which are rectangular areas of website content, including menus, information displays and links.



There's nothing to stop you from using the classic editor, if you prefer it.

that there weren't any hidden limitations on the free version that gave us anything to be worried about.

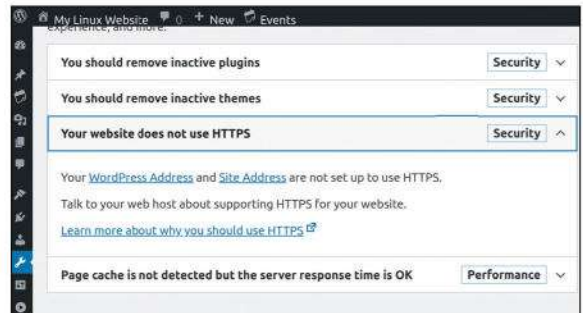
To install a plugin from the search results page, click on the Install Now button and then click on the Activate button when the installation has completed. It's certainly a step up from the old days when you had to manually unzip the plugin and upload it yourself!

Once the plugin is installed and activated, you'll notice the extra features in your content straight away. All Posts and Pages now have social media share icons at the bottom of the page. Unlike the built-in social media icons, the AddToAny icons share the content to the social media page of the reader. By default, AddToAny adds sharing icons for Facebook, Twitter and email, with a plus icon that leads to a huge range of extra social media sites.

Many find these default settings sufficient, but it's worth checking out the settings (Dashboard > Settings > AddToAny), where you can make refinements. In your case, it might be that some of the social media sites are more important than others and deserve to be put on display. You can also make refinements to the placement of the buttons. This, along with some further refinements of the colour and size of the icons, might be needed if there is a clash with the colour and layout of your *WordPress* site.

Add a to-do list

A to-do list can be used to prod fellow contributors or even as a reminder to yourself for what needs to be



The Site Health page is telling us that we don't yet have a valid TLS certificate installed on the site; so we can't use HTTPS.

done. Dashboard To-Do List is a plugin that only had six ratings at the time of writing, but it has over a thousand installations and a 5/5 overall score, so we decided to take a chance on it.

It's a fairly simple plugin. You can add to-dos within the Dashboard with a simple one-line-per-entry interface. This to-do list can be made to appear as part of the front page and entries can be directly added from there. It's just the thing to keep a project moving.

Upgrade the security

If your *WordPress* website will be publicly available on the internet via a domain name, you probably want to enable HTTPS (HyperText Transport Protocol Secure). This establishes an encrypted connection between the browser and the server to ensure that no one else can pretend to be you when it comes to publishing content. Not having HTTPS enabled for your site causes some of the popular browsers to put a line through the URL, to indicate that your site is not secure, and it can affect your Google search ranking.

Fortunately, it's fairly easy to add HTTPS support to your site as it's built into *WordPress*. If you are using commercial hosting for the site, you might not have to do anything as the hosting company should have taken care of providing a valid TLS (Transport Layer Security) certificate. You can find out what the current status is in that regard by examining the site health (Dashboard > Tools > Site Health). If this tells you there is a valid site certificate, you merely have to change the WordPress Address and the Site Address (Dashboard > Settings) so they begin with **https://** rather than **http://**. Make sure the Site Health page really does say that it can find a valid certificate before changing the WordPress Address and Site Address details, otherwise you might not be able to access your site.

If you are hosting *WordPress* on your own server, adding a valid TLS certificate is fairly easy, too. We'll make use of a service called Let's Encrypt (**<https://letsencrypt.org>**) that issues free TLS certificates and a Python script that automates the process. Make sure you are logged into the command line of your server (locally or via SSH). We'll assume you're running a Debian-derived distribution such as Ubuntu. Type **apt-get install python3-certbot-apache -y** to install the utility, and type **certbot --apache -d [domain name of your site]** to fetch and install the certificate.

Return to the *WordPress* Site Health page to make sure it can find the certificate, then follow the same procedure that we outlined above for professionally hosted sites. As before, don't try changing the

» PLAYING AWAY FROM HOME

One of the great things about using *WordPress* as a content management system is that you can log in from anywhere. Writing content on a cloud-based word processor, such as Google Docs, then pasting it into *WordPress* later is something that can work. The two systems are compatible in that *WordPress*'s Gutenberg block editor respects Google Docs' paragraphing, headers and formatting.

The official *WordPress* app for iOS and Android is free and well worth considering. It has an editor for creating Posts and Pages. Beyond this, it also has facilities for the management of content and comments, a fast and convenient way for doing a quick check of what's happening on your site while you're on the move. A limitation of the app is that it doesn't have the full range of site configuration options, but it does contain a link to jump out of the app and into the web browser of your device. This means that you have to sign into *WordPress* again. From there, you are using the mobile version of the Dashboard, which is still usable, although not as neat as the app.



The official Android app is pretty spiffy for content creation and management on the go.

Plugin	Description
<input type="checkbox"/> AddToAny Share Buttons Settings Deactivate	Share buttons for your pages including AddToAny's universal sharing button, Facebook, Mastodon, LinkedIn, Pinterest, WhatsApp and many more. Version 1.8.9 By AddToAny View details
<input type="checkbox"/> Akismet Anti-spam: Spam Protection Activate Delete	Used by millions, Akismet is quite possibly the best way in the world to protect your blog from spam. Akismet Anti-spam keeps your site protected even while you sleep. To get started: activate the Akismet plugin and then go to your Akismet Settings page to set up your API key. Version 5.3 By Automattic - Anti-spam Team View details
<input type="checkbox"/> Classic Editor Deactivate Settings	Enables the WordPress Classic Editor and the old-style Edit Post screen with TinyMCE, Meta Boxes, etc. Supports the older plugins that extend this screen. Version 1.6.3 By WordPress Contributors View details

From the Plugins page, you can add, remove and disable plugins. Click Add New to search for new plugins.

addresses of your site until the Site Health page tells you it can find the certificate or you may end up unable to access your *WordPress* site. It is possible to create a certificate for a local network setup, but it's quite a lot of work and basically unnecessary if the site can't be accessed from the internet.

Add extra users

It could well be that you are the sole user of your website, but for many projects, you need to allow other users to add content and take over some of the management responsibilities. Navigate to the Users page (Dashboard > Users) to begin. This settings page lists the users on the system, and initially, this consists solely of your administrator account. Click Add New to add a new user and you're taken to the user creation page. At the minimum, you must give a new user a username and assign a role. Extra things you can add, such as the full name of the user and a link to that user's website, show up as part of that user's profile. *WordPress* offers a set of predefined roles for a user.

The most useful roles for most projects are Admin, Author and Contributor. As the name suggests, an Admin can handle a lot of the admin and configuration for the site; it's often handy if someone else can take over those duties in an emergency. An Admin can also create and publish posts, and edit and delete other people's posts. An Author can create their own posts and publish them. A Contributor is a bit more limited in that they can create posts, but they can't publish them.

Manage events and dates

As the name suggests, the Events Calendar adds extensive calendar features to *WordPress*. It's ideal for adding events for websites that are based around clubs, for example. The plugin itself is free, but it has a pro version with some extra features and a mixture of free and paid-for add-ons.

Install it like any other plugin (Dashboard > Plugins > Add New), and once activated, you have extra options under the Events heading on the Dashboard sidebar. As with much of *WordPress*, the interface

should be fairly familiar, and the reason for this is that, behind the scenes, this plugin defines a custom type of Page. If you are using the Gutenberg block editor, enable it in the settings (Dashboard > Events > Settings > Activate Block Editor for Events).

Let's create a simple event as an example. Go to the Events section (Dashboard > Events) and click on Add New, as when creating a Page or a Post. This opens up the Page editor, where you can add the title, body text and date and time for the event. As we're using the block editor, you edit these by clicking on them. For example, if you click on the date, a calendar pops up. Usefully, this allows you to select All Day for the time.

As with regular types of Post, click on the Publish button to add the event to the site. It's left to you to add a link to the events page. One way of doing this is to add a link in the navigation bar that defaults to being at the top of the front page of the website. Go to the site editor (Dashboard > Appearance > Editor), and click on Navigation to bring up the navigation block in an editable state. Add a new link by clicking the + button. For the URL, use `[URL of website]/index.php/event/`, and give it a label such as Upcoming Events.

Over to you!

We've covered the basics of setting up a *WordPress* site, along with some typical setup scenarios. It's a system that has been through some big changes over the last few years. The Gutenberg block editor has a lot going for it because it means that nearly all aspects of site design and content creation use a common interface, and it's an editor that allows for fast entry of good-looking content. Plugins can create custom post types (such as calendar entries), further enhancing the consistency for the administrator of the site, and this is another advantage of the way that *WordPress* has changed in recent years.

If you don't, for some reason, get on with the Gutenberg editor, there's nothing stopping you from going back to the classic editor through the use of a plugin. As long as it's maintained by the *WordPress* developers, you're not missing out and you still benefit from most of the newer features beneath the hood. **LXF**

QUICK TIP

We'd recommend switching on the Akismet plugin once you make your website available on the internet. It protects your site from spam comments. It's installed by default but disabled.

The Events Calendar enables you to add events via an interface that's consistent with other post types. These are added to the Events page.

<input type="checkbox"/> Title	Author	Event Categories	Tags	<input type="checkbox"/>	Start Date	End Date
<input type="checkbox"/> Exciting future event	Mike		—	—	November 24	November 24
<input type="checkbox"/> Yet another event	Mike		—	—	November 1	November 1
<input type="checkbox"/> November Picnic	Mike		—	—	October 29	October 29
<input type="checkbox"/> Linux Meeting! — Draft	Mike		—	—	October 28, 2023	October 28

» **IMPROVE YOUR LINUX SKILLS** Subscribe now at <http://bit.ly/LinuxFormat>

BACK ISSUES » MISSED ONE?

ISSUE 309 December 2023

Product code:
LXFDB0309



In the magazine

Manage your apps by harnessing the power of *Podman*, and discover how random numbers are used by the kernel to keep your data safe. Plus, find out how to master printing in Linux, customise your *WordPress* website, tweak your kernel scheduler, use Python to scrape the web, process your smart-home data, and much more. And make sure you read our *Roundup* of Raspberry Pi killers...

ISSUE 308 November 2023

Product code:
LXFDB0308



In the magazine

Bring together all your movie, TV, music, audiobook and photo streaming needs under one roof. Plus, we take an in-depth look at the Pop!_OS distro, explain how networking works, and guide you through tutorials on *WordPress*, fast file management, reviving the BBC Domesday Project and more. And don't forget to check out our packed reviews section, plus we examine the brand-new Raspberry Pi 5.

ISSUE 307 October 2023

Product code:
LXFDB0307




In the magazine

Power up your Pi with awesome AI projects, and take control of your phone with the privacy-focused /e/OS. Plus, find out which open source personal cloud server is best for you, read all about the forthcoming Wi-Fi 7, and discover how to transform images into GIFs via the terminal, manage a VPS, boost your Steam Deck, and much more. And there's our usual news, reviews and sysadmin tips.

ISSUE 306 September 2023

Product code:
LXFDB0306



In the magazine

Discover how AI can (and can't) help with your coding, plus add Pi-style GPIO to your PC, access your home server from anywhere, and write your own text adventure and *Breakout*-style games. Also, *Adminsteria* returns with news and reviews for sysadmin types, we take an in-depth look at how Azure stacks up against AWS, there's a *Roundup* of media-creation distros, and plenty more...

ISSUE 305 Summer 2023

Product code:
LXFDB0305



In the magazine

Don't let Big Tech be in control of your documents – find out how to run your own office apps in the cloud. Plus, learn how Linux certification can kick-start your career, step back in time with our look at the evolution of coding, and read our recommendations for the best open source CAD packages, not forgetting our usual packed tutorial, review and Raspberry Pi sections.

ISSUE 304 August 2023

Product code:
LXFDB0304



In the magazine

Run your own AI and take over your world with local open source machine learning, discover the diverse delights of BlendOS, and learn all about AMD, Intel and Apple's different takes on CPU design. Plus, we have an in-depth interview with Collabora's Michael Meeks, a *Roundup* of official Ubuntu spins, and tutorials on everything from restoring old photos to coding a text adventure.

To order, visit **www.magazinesdirect.com**

Select **Single Issues** from the tab menu, then select **Linux Format**.

Or call the back issues hotline on **0330 333 1113**
or **+44 (0)330 333 1113** for overseas orders.

Quote the product code shown above and have your credit or debit card details ready.

USA? EU? THE MOON*?

UK readers
turn to
p16

SUBSCRIBE!

Don't wait for the latest issue to reach your local store – subscribe today and let *Linux Format* fly straight to you. Faster, cheaper and with DRM-free archive access!



» USA
From \$135.49
For 13 issues

» REST OF THE WORLD
From \$135.49
For 13 issues

» EUROPE
From €119.49
For 13 issues

IT'S EASY TO SUBSCRIBE!

Visit www.magazinesdirect.com/linux-format

Call +44 0330 333 1113

Lines open Monday-Friday, 9am-5pm, UK time.

*We don't actually deliver to the Moon. Yet.

BBC BASIC

Credit: www.bbcbasic.co.uk/bbcSDL

Relive your BBC Basic days!

Turn your home system into a BBC Micro by following **David Bolton's** advice to download BBC Basic and write programs just like it's 1982.




**OUR
EXPERT**

David Bolton learnt to program in Basic back in the dark ages, which helped him get an O Level in computer science long before personal computers were even a thing.

The BBC Micro was launched 42 years ago, in 1981. Like other home computers of that era, such as the TRS-80 and CBM Vic-20, it came with a built-in Basic interpreter. Unlike most other home computers, this wasn't created by Microsoft but was one that Acorn had created.

A former BBC engineer, Richard T Russell, who was involved with the BBC Micro development, has developed an excellent open source version of BBC Basic. You can find the source code on GitHub (<https://github.com/rtrussell/BBCSDL>). It's cross-platform and comes in two versions: BBC Basic for SDL 2.0 and BBC Basic Console Mode edition. It's available on platforms such as Linux, Windows, Android and Raspberry Pi. We're using a Pi but you can follow along on PCs, too.

After you've installed it, take a look at some of the example games and programs that come with it. All include source code and are pretty good. If you've ever wondered how to program a sudoku game including the generator and solver, it's all there. The *Aliens* demo



The two IDEs that you are offered when you start BBC Basic.

shows moving rotating objects at a high frame rate, achieving 60fps for 80 objects on a Raspberry Pi 4B.

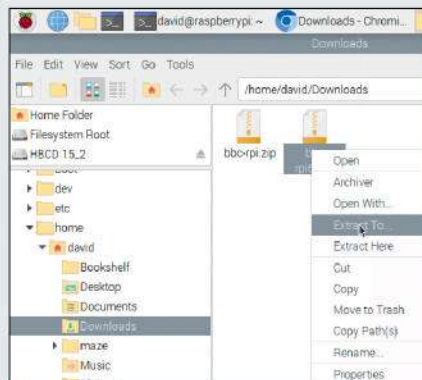
The two IDES

When you click the BBC Basic icon, you are offered two different IDEs: *BBCedit* or *SDLIDE*. After trying both, if you decide you want to stay with your favourite, just tick the Don't Ask Me Again checkbox. However,

» INSTALLING BBC BASIC SDL


There are two versions of BBC Basic, the console version and the SDL version, which has better graphics. In a browser visit www.bbcbasic.co.uk/bbcSDL. Scroll about halfway down the page and you'll see a table with downloads for different platforms. Select the one appropriate for your system, usually 64-bit x86; we're going with 64-bit Raspberry Pi. If you're not sure, use the terminal command `uname -m` and it should show `aarch64` on 64-bit Pis (x86-64 for PCs) or `armv7l` for 32-bit Pis (x86 on 32-bit PCs).

After the ZIP file has downloaded, in File Manager open the **Downloads** folder, then right-click and on the pop-up menu click **Extract To**. Pick a folder (your **home** folder is fine) and it creates a **bbc-rpi64** folder (on 64-bit or **bbc-rpi** on 32-bit). In that folder, you should see an



File Manager showing the Extract To menu item for the downloaded ZIP file.

examples folder, a **lib** folder and some other files. Find the **install-rpi.txt** file and right-click to open it in a text editor. Now copy the six lines that start `sudo apt-get`



Default warning popup on Raspberry Pi when you try to run an executable. It can be turned off!

install including the two `chmods` and paste them into a terminal, then press Enter. You may find it already has some of the SDL files installed, but don't worry.

Right-click **Open** on the **mkicon.sh** file to create a desktop icon. Double-click it to start BBC Basic. It asks if you want to run an executable each time. You can disable this pop-up by clicking on the File Manager and in **Edit/Preferences**, on the **General** tab, click the checkbox that says **Don't Ask Options On Launch Executable File**.

we recommend you leave it unticked as both have their uses.

SDLIDE was written by Richard T Russell and has a more modern look. It's the better of the two because it includes an explicit Debug Program command and when you run that it opens a window showing all variables with their current values. The other IDE (**BBCedit** by Andy Parkes) seems to offer more features apart from debugging, but it's a matter of individual taste; it does have a Recent menu, which is handy for revisiting files you opened earlier. Both IDEs are a single file at a time, which is a bit limiting when you're used to editors like **VS Code** that can have multiple files opened. However, you can run multiple iterations of **SDLIDE** and **BBCedit**, and have them all open at the same time, and use one for recent files and debugging on the other.

Press **F1** on **SDLIDE** and it opens up the BBC Basic for SDL website. On **BBCedit**, it opens a help window listing the commands.

Graphics on the BBC Basic

There are two main commands for setting up graphics. The **mode** command sets the screen size and pixels. Modes 0 to 7 are the same as the original BBC Micro but there are now modes 8-33. Mode 7 gives you teletext. Included in the examples under General is a Ceefax demo in 1,134 lines of code. It looks good and works, so if you want a bit of '90s nostalgia, fire it up.

The second command for graphics and characters is **VDU**. BBC Basic has an emulator for these commands – www.rtrussell.co.uk/bbcwin/manual/bbcwin8.html lists all the codes used with **VDU**. When the BBC Micro came out in 1981, monitors were often called visual display units (VDUs), hence the name.

The best way to demonstrate the graphics is to write a small program, in this case a simple tennis game played with paper and two dice. This game was originally published as a paper and pencil game in a UK magazine *Games & Puzzles* back in the mid-'70s.

Understanding BBC Basic

Before we get on to the tennis game, it's important to understand the type definitions at the end of variables.

- 32-bit integers: single % at the end, such as **count%**
- 64-bit integers: %% at the end – **background%%**
- Byte variables: & at the end – **jumpKey&**
- String variables: \$ at the end – **name\$**
- System variables start with an @ – **@platform%**

These are predefined.

These also apply when you have arrays, such as **Dim names\$(10)**

The game was originally developed in a graphics mode using code that the examples provided, but it was blighted with no room and dim space errors, and switched to a text mode game instead in mode 12.

Anyone for tennis?

The game was originally designed to be played on a sheet of paper with 12 rows numbered 1-12. The net is between the sixth and seventh rows. The players each



Screenshot of the Ceefax program. It's a working program that fetches data from the web.

roll a dice and the player with the highest roll serves. If they are equal, roll again.

The ball, which can be a coin, starts just off the board on the 0th row. The serving player rolls two dice and the coin is moved that many rows if the dice total seven or above. Less than that is a fault and the non-serving player gets the point.

Assuming the coin clears the net, the non-serving player now gets to choose to roll one dice or two, and moves the coin towards the serving player by the total. If it fails to reach the net or goes off the far side of court, the server gets the point. Scoring is as in normal tennis: love, 15, 30, 40, advantage and so on.

The program lets you play against a computer opponent. It has to draw the tennis court, then move the ball by the number of spaces, pick moves and let you choose one or two dice each time. On top of this, it has to keep the score.

The source code for the tennis game is on GitHub (<https://bit.ly/lxf310tennis>). Just click the View Raw link and it downloads. Note: you should install BBC Basic to view it, as .bbc files are tokenised. You can save the source code out as a text file by selecting Text File on the File Save As dialog in **SDLIDE**.

Court by surprise

BBC Basic uses special characters to set colours.

This line in **procDrawCourt** outputs a string in green. There's an array of 12 strings, **Row\$(12)**, which holds the board. This is initialised in the **proclnitCourt** call.

```
Row$(row%)=chr$(17+chr$(128+2)+ROW$(row%)+chr$(17+chr$128
```

You can see this type of thing in the Hanoi game, where lots of different coloured disks are drawn. The 2 in **chr\$(128+2)** sets the colour. Each bar includes the row number in text. This line after the **for** statement uses **row%** (an integer that counts from 1 to 12).

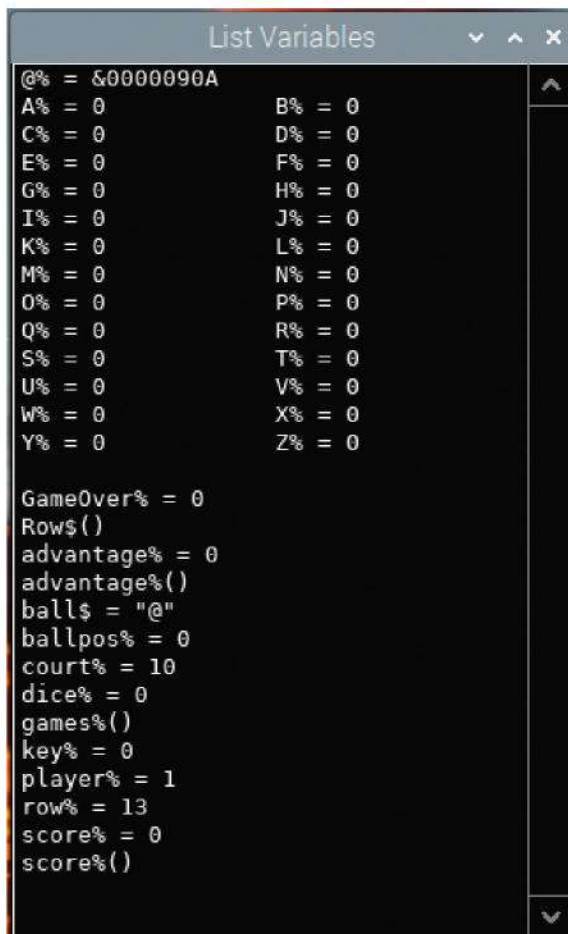
```
Row $(row%)=string$(25+(row%>9)," ")
+str$row%+string$(25," ")
```

To make sure the lines remain the same length when the row number goes from 9 to 10, the **row%>9** expression evaluates to true and in BBC Basic, true has a value of -1 and false is 0. So, for rows 10 through 12,

QUICK TIP

SDLIDE is the IDE you want if you are doing any debugging. It also gives you the option of having all BBC Basic keywords in upper case for that shouty '80s Basic look or lower case for a more modern approach.





The debug variables window opens when you start debugging.

QUICK TIP

Occasionally, the www.bbcbasic.co.uk site hits a bandwidth exceeded limit. If you search for a link in Google and it takes you to a bandwidth exceeded page, go back to the search and click on the three dots after the link in the search results. This pops up a window and it may have a cached button. Click this for a copy of the page.

there are 24 spaces before the number and 25 for the other rows.

The gameplay screen (see *opposite, top*) won't win any prizes for appearance. The court is split so the net appears in the middle. The @ is the ball in the serving position and it moves down as you press 1 or 2.

The **print tab(x,y)** is used to position the point of printing at column x on row y, used for drawing all the text including the scores and fail messages.

The program is designed so the main logic is in the first half and the rest is all the procs and functions.

All variables are either strings or integers. This is the code for drawing the court, net and scores:

```
def procDrawCourt
cls
print tab(40,court%-3)"Score Player ",fnScore(score
%(1,advantage%(1))" computer "fnScore(score%(2),adv
antage%(2))
for row%=1 to 6 step 1:
print tab(5,row%+court%)Row$(row%);
next
print tab(5,7+court%)string$(51,"_")
for row%=7 to 12 step 1:
print tab(5,row%+court%+1)Row$(row%);
next
print tab(27,ballpos%+court%)ball$
endproc
```

The game loop

This starts around line 18, the **while GameOver%=0**. After reading the key, which is limited to 1, 2 or Escape,

the variable **dice%** is set to 1 or 2 accordingly and used to randomly set how far the ball goes, calculated in **fnMoveBall**. The position of the ball is held in **ballpos%** and that must be in the range 1-12 or the ball is out. It also has to be in the opponent's side of the net or you lose a point. The function **fnBallOut** has the logic to work out if the ball is out of court.

The computer player just uses a simple algorithm to decide between one and two dice in the function **fnComputerMove**. If it's nearer the net – that is, **ballpos% < 10** – then it's one dice, otherwise it's two. It would be a fun exercise to have two computer players play a few thousand matches and try different one/two dice-deciding algorithms to see which plays best.

Some programming tips

BBC Basic for SDL isn't compiled and only does error checking when the program is run. This can lead to a frustrating search to find what has gone wrong. One disadvantage of not having any line numbers is that the On Error trapping can't tell you what the offending line number is. Late in the day, though, it turned out that there is a way to do it.

First setup a **procError** procedure. This is called by an **On Error** line, which should be the first line of your program. Something like this:

```
on error procError : end
```

The **procError** can be fairly simple, like this:

```
def procError
print "Error no ";err; " on ";erl
report
endproc
```

The keywords **err**, **erl** and **report** are built-in error functions and procedures. If, for example, you have a syntax error on a line, you'll get something like 'Error no 16 on 0' and 'Syntax Error'.

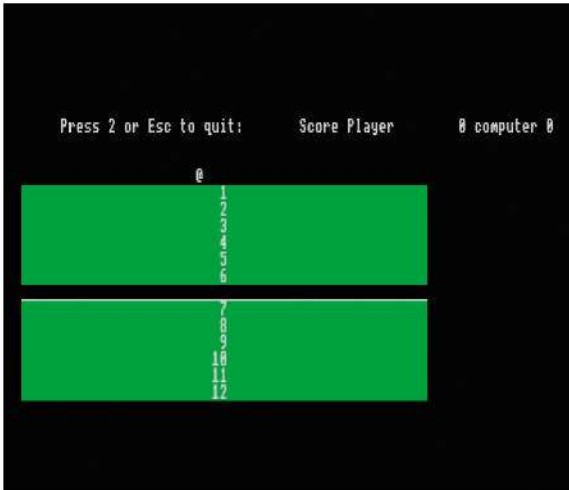
The **0** from **erl** is because there are no line numbers. However, **SDLIDE** has a **Renumber** command on the Utilities menu. This adds line numbers to every line and you can specify the start number and interval. It defaults to starting at 10 and going up by 10, but you can change that.

With line numbers added, now when you run the program, you'll get the line number in the error message. Then, once you've finished debugging it, use the Utilities > **Renumber** command again and this time tick the checkbox at the bottom of the pop-up window labelled **Remove Unused Line Numbers**. Some people find that line numbers clutter up a listing, so prefer to work without them. Note: if you have old programs with **Goto** and **Gosub**, you have to keep those line numbers in. The **Renumber** command strips out the rest, though.

There's nothing wrong with passing a parameter into a proc or function, but if it has the same name as a globally declared variable, things seemed to get a bit confused. Possibly it's a bug, but it's bad programming anyway. Changing the name of the parameter fixed it, so if you get odd bugs, look for those.

Watch your indentation

Many structures need a closing tag, such as **endproc** for **def proc**, **endwhile** for **while** and **endif** for multiline **if then**. If you have code that appears indented wrongly, it suggests you've left off an end tag. Functions don't



Here is the tennis game at the start of a game.

use an **endproc** or **endfn**. Instead, you just assign a value with **=value** and that's it.

Debugging

SDLIDE supports debugging in a couple of different ways. You can sprinkle breakpoints on most lines and when you click the Run-Debug command, the program stops there. Or you can step line by line with F5/F7, or run up to a line with the F6 button and stop there. Most useful is the Debug Variables window that opens automatically when you start debugging. This lets you see all the static variables and all your global variables. The top two columns are the built-in static variables that keep their values between runs. Below them are the global variables from the program.

At one point we had a **player** variable as well as a **player%**. It was an error and it can easily happen. Both variables showed up in the debug window. It's easily overlooked in a source listing, but if you see it here, you know that it exists and can track it down.

The screenshot (*below-right*) shows the main window when debugging. The blue square on the left is a breakpoint and the blue highlighted code is where the debugger is currently looking.

Garbled graphics

BBC Basic is very well done and, quite remarkably, the programs look the same whether running on Raspberry Pi, Ubuntu or even Windows, other than the title bar, which reflects each operating system.

There is quite a learning curve, because there are a lot of keywords, but the online documentation is very extensive. It can't compete for speed of execution compared to a fully compiled program, but it is probably one of the fastest; the *Aliens* demo mentioned earlier demonstrates that. If you need faster speed, the support for assembly built into BBC Basic makes that possible. You can embed assembly

» WHAT IS BASIC?

In the early '60s, computers were still difficult to program. You had to be a scientist or a mathematician. Basic (Beginners All-purpose Symbolic Instruction Code) was created at Dartmouth College for students in non-scientific fields to be able to write programs.

All the home computers in the '80s came with a Basic interpreter, but the version supplied with BBC Micros was better than the standard 8K Microsoft Basic that others used and it was also better optimised, so programs ran faster.

It feels like BBC Basic's designer looked at Pascal and borrowed a few constructs. For instance, loop structures using **repeat** and **while**, name and structure variables (like **c structs**), plus **procs** and **FNs** (functions), where you can pass variables by value or reference and have local variables. Also, line numbers are optional – you only need them if using **GOTO** or **GOSUB** statements that need a target line number. If you are used to a Basic where every line has to have a number, this is revolutionary. For more about the BBC Micro and BBC Basic, see <https://central.kaserver5.org/Kasoft/Typeset/BBC/Contents.html> and www.bbcbasic.co.uk/bbcwin/tutorial/.

language directly in your Basic programs and the built-in assembler is probably better than most other programming languages.

Although this is the graphic version, attempts to write this as a graphics program were unsuccessful. For some reason, the **gfxlib** just did not want to play nice, although it clearly works in the various demos. It's a personal choice, but be sure to tick the Dark Mode boxes and have syntax colouring and lower-case keywords both ticked on the **SDLIDE** Options menu.

Some people turn their nose up at Basic because in the past it encouraged a very lax style of programming, with **Gotos** all over the place; so-called spaghetti programming. This is a modern Basic and makes well-laid-out and easy-to-understand programs possible using modern features such as **while** and **repeat**.

Note that for screenshots on the Pi, we had to use the previous version of Raspberry Pi OS, as the newest one is based on Wayland and many screenshot utilities, such as *Scrot* and *Lazam*, don't work with Wayland. **LXF**

QUICK TIP

Along with **Repeat** and **While** keywords, **BBC Basic** lets you write entire programs without needing a single line number or **Goto/Gosub**. This aids both readability and program structure.

```

if GameOver%=0 then
  player%=1
  distance%=fnMoveBall(dice%,1)
  ballpos% = ballpos% + distance%
  if fnBallOut=1 then
    procAddScore(2)
    procClear(court%-3)
    print tab(10,court%-3)"Player hit ball out"
    wait 500
    ballpos%=0
  else
    player% = 0
    distance%=fnComputerMove
    ballpos% = ballpos% - distance%
  
```

The blue highlighter is the current debug line with a breakpoint in the left gutter.

» **IMPROVE YOUR LINUX SKILLS** Subscribe now at <http://bit.ly/LinuxFormat>

ADVENTURE GAME STUDIO

Credit: www.adventuregamestudio.co.uk

Make your own point-and-click adventure

Nate Drake invites you to relive the glory days of Scumm-style point-and-click adventure games by creating your very own.



OUR EXPERT

Nate Drake is a journalist specialising in cybersecurity and retro games. The first thing he did when discovering Linux in 2004 was play *Beneath a Steel Sky*. Be vigilant.

The '80s and '90s were the heyday of the point-and-click adventure game. The setting of titles varies but the basic premise remains – that of moving your character around a screen using your mouse, interacting with people and objects, and usually adding items to your inventory.

These aspects of point-and-click games were one of the most maddening for those who grew up in an era when you couldn't simply DuckDuckGo a solution, because these games were deliberately designed to make the player take counterintuitive actions.

Who would have thought, for instance, that in order to defeat the undead villain in *The Secret of Monkey Island* (1990), the protagonist needs to spray him with a bottle of root beer? In the cyberpunk adventure *Beneath a Steel Sky* (1994), the game actually requires you to volunteer the character's testicles to an unscrupulous doctor once all of your other trading options are exhausted.

The fiendish level of difficulty is probably why the point-and-click genre has never been as popular as first-person shooters or racing games. Still, there's no beating the little dopamine rush you get once the pieces click together when you use a particular item in a specific place, and progress in the game.

In this four-part series, we're going to explore the mechanics of setting up a simple point-and-click

adventure of your own. There's sample artwork to play around with, as well as minimal coding.

Point-and-click platforms

Readers old enough to remember the golden era of 2D point-and-click games will recall LucasArts titles such as *Maniac Mansion* (1987) and its successor *Day of the Tentacle* (1993). Many such titles were powered by the SCUMM (Script Creation Utility for Maniac Mansion) game engine, which was designed to help developers create graphical adventure games with ease. SCUMM has its own scripting language and lives on to this day through the open source *ScummVM*, which can reinterpret the script to run on more modern platforms.

In order to write this series of tutorials, we seriously considered SCUMM and a number of other platforms for creating point-and-click adventures before settling on *Adventure Game Studio*, which is supported by more recent *ScummVM* releases.

This open source software was first developed by British programmer Chris Jones in 1997 as *Adventure Creator*, so crucially it's originally from the era of games we're trying to recreate. The team holds annual competitions for the best point-and-click games and the platform's been used to create commercially successful titles. Even more crucially, the game editor is capable of creating games that can be run in Linux.

Getting started

Having praised AGS's ease of use and flexibility, sharp-eyed readers will have noticed a catch: the editor was originally developed for MS-DOS (the platform of the era), so is only available for Windows. In our tests, we were able to get the software running on Ubuntu 23.10 using the *Wine* compatibility layer. Start by opening the terminal and enabling support for 32-bit architectures by running the command:

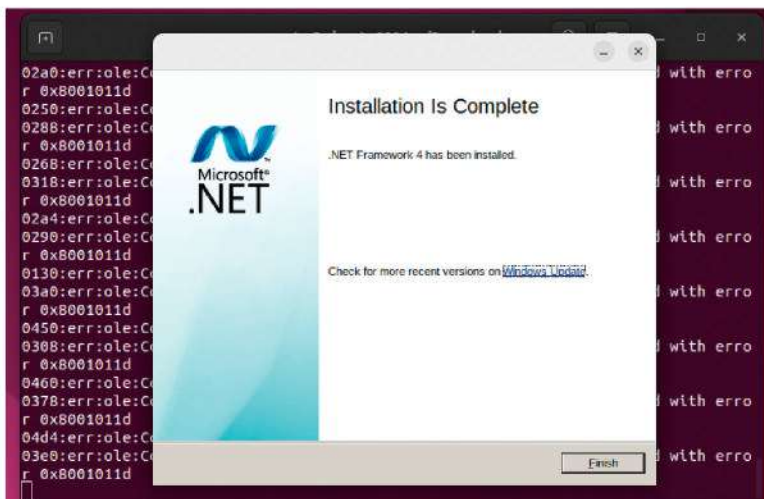
```
$ sudo dpkg --add-architecture i386
```

Next, install *Wine* from your distro's repositories:

```
$ sudo apt-get install wine
```

If you prefer to use the most recent version of *Wine*, you can manually add the repository and download directly by following the instructions at <https://wiki.winehq.org/Download>. Next, open your browser and head to www.adventuregamestudio.co.uk/site/ags/.

AGS games run on Linux but the editor is Windows only. You should be able to run it using Wine.



Click to download the latest full version of AGS. Return to the terminal and go to your **Downloads** folder with:

```
$ cd Downloads
```

Next, run the installer using `wine` – for example:

```
$ wine AGS-3.6.0.53-P6.exe
```

Setup should launch and inform you that you need to install at least version 4.6 of the .NET framework. Click OK to launch the Microsoft website within your browser. We found the best results by downloading .NET Framework 4.6.2 Runtime. Once downloaded, return to the terminal and run the installer:

```
$ wine ndp462-kb3151802-web.exe
```

Click Next, then Install. Once install is complete, go to the terminal and use `wine` to run AGS setup again.

Note that while these steps worked perfectly in Ubuntu 23.10, when we tried to do the same in a clean install of Ubuntu 22.04 LTS, setup failed because AGS couldn't launch the configuration files. You may be able to fix this by installing the package **winetricks** and installing the necessary .NET framework with:

```
$ winetricks dotnet40
```


Even if you do get it running, the *Wine* database shows AGS has some compatibility issues, such as being unable to run at certain resolutions. If you have trouble, try installing AGS in a virtual machine running Windows. Any games you create will run in Linux.

AGS setup

Whether using Wine or Windows, AGS setup is simple. Click Next to accept the licence agreement, confirm your install location, then click Next again.

The default installation method is Full, which you should leave as is, given that it includes support for creating Linux games. Click Next to go to the additional tasks screen. This includes options such as installing support for C++ and creating a desktop icon. We suggest leaving these ticked. Click Next and Install.

If all goes well, you'll see the option to launch AGS immediately. Click Finish to continue. On the welcome screen, click Continue to launch the New Game Wizard. You're first asked to choose a template on which the



Edges define exactly how far the character can move before a new 'room' loads. Click and drag to change them.

game is based. Unless you know what you're doing, we recommend selecting BASS.

This provides an interface similar to *Beneath a Steel Sky*, whereby your character can walk around, interact with objects and use an inventory at the top of the screen. Most importantly, it comes with some default graphics to use until you're comfortable using AGS.

Give your game a name, such as Nate's Adventure, and enter a filename for the game itself (choose wisely, as you can't edit this later). By default, the game folders are located in **documents**, but click the three dots to change or create a new folder, then hit Finish.

Main window

AGS now launches the main window. You're going to be spending a lot of time here, so it's worth familiarising yourself with the interface. At the top-right, you'll see your Project Tree. This enables you to explore various aspects of your game, such as colours, sprites, characters and so on. When you click in any of these, their properties appear at the box in the bottom-right.

Tabs for any currently open windows are displayed along the top of the main window. In order to get

QUICK TIP

If you do choose to change the background for Room 1, it's best to ensure the image is the same resolution as that for the game you specified in General Settings.

» KICKING IT OLD SCHOOL: USE 256 COLOURS

If you've read our warnings about using 256 colours but are bent on reliving the glory days of 8-bit gaming, you may want to rebalance the palette. Click on Colours in the Property Tree, then on the Palette tab in the main window.

Your GUI and main character graphics use the gamewide colours displayed. Those marked with an X are reserved for background colours, which differ from room to room.

By default, this severely limits the number of colours available. To fix this, select the first X slot. Hold Shift and then click on the last slot in the row marked '80'.

Look to the Properties box at the bottom-right, then change ColourType from Background to Gamewide, and

you'll see the selected Xs turn into various colours.

If you select one of these colours, you see its RGB values in the property grid, and can adjust the colour of that slot. If you don't want to do it manually for every colour, set up slots 42-95 in a paint package, save a BMP file, right-click on the selected slots, and choose Replace Selected Slots From File.

It's best to set your palette exactly how you want it now, because if you change it after adding graphics, you'll probably have to reimport them.

You can find more help with this in the dedicated AGS Colour Editor tutorial available at: <https://adventuregamestudio.github.io/ags-manual/ColoursEditor.html>.



If you insist on having an 8-bit game, take the time to add as varied a colour palette as possible.



QUICK TIP

The Ubuntu repos do contain a program named AGS but this is only the game engine. In order to make use of it, you first need to create a game in *Adventure Game Studio*. You can then run the game using the AGS engine from your Linux OS.

started, scroll to the very top of the Project Tree pane and click into General Settings.

These are global settings that apply to your entire game. Scroll down to Characters Turn Before Walking and click to change this to True. This means that if the character is facing right, for instance, and you click to move him left, they turn around first, which makes the game movement more realistic.

Picking your palette

If you scroll to the very top of General Settings, you'll notice the Basic Properties section contains a Colour Depth value. This is where you can choose to create a palette based on 8-bit, 256-colour or hi-colour.

Unless you know what you're doing, we suggest you leave the 32-bit setting as is; 256-colour games often don't play well with modern graphics cards, as they aren't supported by OpenGL or Direct3D drivers.

This said, if you're determined to make this change, now's the time, as any graphics you import are set to your chosen colour depth. In other words, if you change the colour depth halfway through the game, you need to reimport all your graphics.

Basic Properties also enables you to set the game resolution (default is 320x200). Higher resolutions allow more detailed graphics, but the game can run more slowly. If you're unsure, leave this setting as is.

A room with a view

If you're serious about creating a point-and-click adventure, you probably know that they consist of a series of 'rooms'. The name is a misnomer as rooms can be outdoor areas; it simply represents a defined area with its own background and interactive objects.

If you look at the Project Tree on the right of the screen, you'll see a Rooms category. Click the + to expand this, then do the same for Room :1 listed there. Next click Edit Room. You'll see that there's already a science-fiction-style background for this first room, kindly contributed by AGS Forum champion Selmiak. You can load a new background later on if you

wish by clicking the Change button, but just leave it as it is for now.

For this next part, feel free to use the slider at the top-right to zoom out a little, as we'll be working with the borders of the image. When this window loads, you'll notice the Room Parts section. There's a box marked Edit This Room's... – click the three dots beside the word Room, then choose Edges.

You should now see four yellow lines drawn around the background image. These lines represent exactly how far your character should walk for the game to decide when they have left the room. Experiment with dragging these a little to define the boundaries of this room before going any further.

When you clicked on Room before, you'll have seen a padlock icon next to the various options. You can click this once you're satisfied with edits you've made to stop yourself from making any accidental changes. You can also click on the eyeball icon next to the padlock to show or hide features such as Edges.

Walking on by

By default, your character can walk anywhere on screen. For this default background, this doesn't make much sense, as a lot of it is taken up by outer space.

This is why we need to define walkable areas – places within which your character can walk without doing something illogical, like floating on air.

Return to the Room Parts section and click the three dots beside Room > Walkable Areas to select Walkable Area ID 1. This again is an example the game has already loaded for you, showing an area (coloured in blue) where your character can walk.

As you can see, the character currently can't go any further than the right-hand edge of the leafy area. Still, it would be good to allow him to walk down the grassy verge to the leafy area at the bottom-right, so let's add a new walkable area, overlapping the one provided.

To do this, click on the three dots to the left of Walkable Area ID 1 and choose Walkable Area ID 2. This informs AGS we're going to define a new walkable

» WALK ON BEHIND

If you've followed the main tutorial, you already know how to set walkable areas. AGS also supports walk-behind areas. These are sometimes called priorities in other game editors. Quite simply, they're an area the character needs to walk behind rather than in front of.

Setting these up is very similar to walkable areas. Just open the Room tab, and in Room Sections, choose Room ... Walk-Behinds. You can then use the drawing tools at the top of the window to define the walk-behind area.

For instance, when researching this tutorial, we used the rectangle tool to make the monolith containing the glowing orb a walk-behind area.

For this to work properly, you need to set a baseline. This tells the game

where the character has to be in order to be drawn behind the area.

This is normally placed at the lowest point of the walk-behind area. Move your cursor to the bottom of the walk-behind area (in this case, the rectangle we drew) and look at the Mouse Position display above the room background.

These are the X and Y coordinates of your mouse cursor (displayed at the top-left of the window). Type the Y coordinate in the Baseline field in the Properties box – 143, for example. Press



Walk-behind areas are defined the same way as walkable areas. Play around with this feature in the example room.

F5 to save and run the game. Try to walk through the walk-behind area. If it doesn't work, it's most likely that you need to tinker with the baseline, so double-check your Y coordinate.

area. Take a moment now to look at the new buttons that have appeared along the top of the screen.

These are basic drawing tools of the kind you'd find in a simple image editor. You can use these to mark out the new walking area. Before starting, make life easier by increasing the zoom, then scroll down and across to the area of the background on which we'll be drawing.

The official AGS documentation says that games check these areas against the bottom-centre of your character. This means it's wise to create the walkable area a little thinner than necessary, to allow room for your character's feet to spill out on either side.

During our tests, we found the easiest way to define a walkable area is to use the Freehand tool to create a rough loop. You can then use the Fill Area tool to fill it in completely. Alternatively, use the Line tool to create a neat series of interconnected lines to fill in. You can use right-click to erase as you go. Just make sure your new area (highlighted in green) overlaps with the other walkable area (currently highlighted in grey), so the character can move freely from one place to another.

Once you've finished defining this new walkable area, it would be a good idea to redefine the bottom edge of this room so that when your character walks down this leafy outcrop, they can go to a new area.

To do this, click the three dots beside Room in Room Sections and choose Edges again. You can now raise the bottom edge to intersect with this new area.

Click the save icon or hit Ctrl+S when you're done. If you want to test out your new walkable area, hit F5 to run the game at this stage. Your default character, Roger, should now be able to walk to the outcrop at the very bottom-right of the screen. At this stage, nothing happens because we've yet to define any extra rooms. Click the X at the top-right when you're done.

Crafting your character

Since we haven't yet defined a player character, it may have come as a shock to see Roger fully formed and walking around when you ran the game. This is another example thoughtfully provided by AGS.

If you want to see the images that go into Roger's makeup, click Sprites in the Project Tree, then Defaults on the left. You can right-click individual sprites to edit them in your default editor (they're in PNG format). You can also click Import New Sprites From Files.

We'll return to this screen later in the series, but for now, we're interested in getting the mechanics of the game working, so don't worry too much if Roger doesn't resemble your ideal player character.

If you wish, you can expand Characters in the Project Tree, then click on Roger to change his name in the Properties window. Make sure you change both the RealName – Nate, for example – and the ScriptName – for instance, cNate – to one of your choosing.

Warm up your hotspots

One of the staples of any point-and-click game are the hotspots. These are different from objects that you can pick up and add to your inventory.

Hotspots are essentially an area of the background that a player can look at and interact with. They're



usually used for fixed items that won't change over the course of the game, such as trees or buildings.

Once again, AGS has made life easy for you by already inserting a hotspot into the game for you in the form of the glowing orb inside the strange stone structure. To get to grips with this, in Room Sections, click the three dots next to Room and choose Hotspots. Click the dots again to choose **hGlowingOrb**.

The orb now appears highlighted in blue. In the Properties box at the bottom-right, the Description field is set to Glowing Orb, which is what appears during the game when you hover the mouse over it.

Currently, if the player right-clicks to Look At the orb, nothing happens, so let's make the game more interesting by adding a more detailed description. In the Properties box, click the lightning bolt icon to open Events. Select the field named Look At Hotspot, then click the three dots in the empty field next to it. This launches the in-game script editor.

In between the curly brackets, insert:

```
Display("It's a glowing orb made of some strange metal.");
```

Next, it's time to add your very own hotspot. Return to the Room1 tab along the top of the screen. Click the three dots beside Hotspots and choose **hHotspot2**. Use the drawing tools along the top to mark out the blue spaceship on the far-left of the screen. In the Properties box, bottom-right, click on the Properties icon. From here, you can set a proper description (Spaceship) and name (hSpaceship) for this hotspot.

Next, click the Events icon and again click the three dots beside Look At Hotspot. You can now use the script editor to display text for the spaceship when the player right-clicks on it, for example:

```
Display("An alien blue craft. If only I knew how to fly it...");
```

Now press F5 to save and run the game. Right-click on both hotspots to check the descriptions display OK.

To be continued...

By now, you've built the first room of your point-and-click adventure with clearly defined edges. You have also set walkable areas and created hotspots, which the player can examine. Stay tuned for next month's tutorial, where we'll cover interacting with objects and managing your inventory. **LF**

Right-click on a hotspot to examine it. You'll learn how to interact with hotspots in other ways later on.

QUICK TIP

If you're using Wine and choose to create a desktop shortcut in Ubuntu, by default this is disabled. Right-click the icon and choose Allow Launching to fix this.

» **TRY OUR OWN CLICK ADVENTURE** Subscribe now at <http://bit.ly/LinuxFormat>

ADMINISTERIA

Massively improve your SSH login security

Passwords are so last century proclaims **Stuart Burns!** Using public keys and a couple of tweaks makes for a far more secure login.

SSH keys can ease admin burdens, so we're going to explain how to configure *OpenSSH* to require SSH keys for login, tightening system security considerably, as well as enforce login via key only (as opposed to allowing keys or passwords) and how to disable root login via SSH (script kiddies love to scan for these and brute force them).

Root login using username and password will still work at the physical console (and out-of-band consoles). If an attacker has access to the console, it's trivial to reboot and reset the password anyhow.

Linux systems administrators can bolster the security of their servers quite easily, guarding from low-level attacks by forcing all SSH logins to use SSH keys for a secure login. To be clear, we are only discussing SSH logins, not app logins and so on.

This article assumes you have a keypair generated already. If not, it's not difficult; just use the following command in a local terminal on your laptop or desktop:

```
$ ssh-keygen -r rsa
```

The `ssh-keygen` command generates the keys in a subdirectory of the user's **home** directory called `.ssh`. Inside this folder are two files: `id_rsa`: (your private key that you must keep secure) and `id_rsa.pub`, your public key that you can distribute and place on other systems.

You can generate a set of keys for each laptop and desktop if desired and copy them across as needed. It just depends on the need and convenience.

So, before we start, what do keys actually bring to the table? Secure Shell (SSH) keys have become the de facto standard for securing remote access to Linux servers. Keys offer several key advantages over traditional password-based authentication for many reasons. Not least because of:

Enhanced security: SSH keys provide a more secure way to authenticate users compared to passwords. Passwords are susceptible to brute-force attacks, dictionary attacks and social engineering. In contrast,

SSH keys use strong cryptographic algorithms, making them extremely difficult to crack.

Bye-bye passwords: One of the most significant benefits of SSH keys is the elimination of password-based logins. This reduces the risk of password-related vulnerabilities such as weak passwords, password sharing and password reuse.

Making life easier: SSH keys offer a seamless login experience for authorised users. There's no need to remember or type complex passwords, making it convenient for administrators and users alike.

Toughens up the login security: With password authentication, there is always a risk of exposing the password during login attempts. SSH keys, on the other hand, are stored securely on the client machine, reducing the risk of interception during transmission.

Implementing SSH keys not only enhances security but also eases the administration burden for Linux system administrators, because there's:

Reduced password management: With SSH keys, administrators no longer need to manage user passwords or worry about password policies. This leads to significant time and resource savings.

Centralised key management: SSH keys can be centrally managed, enabling administrators to control access to servers more efficiently. Key revocation and rotation become straightforward tasks.

Enhanced auditability: SSH key-based authentication provides a clear audit trail of who accessed the server and when. This can be invaluable for compliance and security monitoring.

The really easy way to do that is to use the `ssh-copy-id` command. It does exactly what it says on the tin and the best bit is that it does any required setup for you. Copy the key across like so:

```
$ ssh-copy-id stuart@mymachine_to_login_to
```

Configuration is key

To utilise the benefits of SSH keys and tighten security, Linux sysadmins need to configure *OpenSSH* to require only SSH keys for login. Here's how to do it:

The *OpenSSH* server configuration file is usually located at `/etc/ssh/sshd_config`:

```
$ sudo nano /etc/ssh/sshd_config
```

Search for the **PasswordAuthentication** directive. By default, it is set to **yes**, allowing password-based logins. Modify it to **no** to disable password-based authentication. This enforces the use of SSH keys:

```
PasswordAuthentication no
```

These files are the private and public keys that can be used.

```
stu@workstation: ~/.ssh$ ll
total 24K
drwx----- 2 stu stu 4.0K Oct 20 16:06 ./
drwxr-x--- 37 stu stu 4.0K Oct 22 15:14 ../
-rw----- 1 stu stu 2.6K Oct 22 15:18 id_rsa
-rw-r--r-- 1 stu stu 569 Oct 22 15:18 id_rsa.pub
-rw----- 1 stu stu 2.2K Oct 20 14:02 known_hosts
-rw----- 1 stu stu 1.4K Oct 20 14:02 known_hosts.old
stu@workstation: ~/.ssh$
```


Save the changes and exit the text editor. To apply the new configuration, restart the SSH service:

```
$ sudo systemctl restart ssh
```

As a bit of a pro tip, you can see what is happening during the login from the client using SSH, adding **-vvv** as a parameter. It can give useful feedback if there are issues.

Even more secure

There are additional steps you can take to further tighten the security of your Linux server:

Disable root login: To prevent direct root login via SSH, set the **PermitRootLogin** directive to **no** in the SSH config file. Administrative tasks can be performed using **sudo**:

```
PermitRootLogin no
```

Use strong key pairs: Encourage users to generate strong key pairs with robust passphrases. The use of passphrase-protected keys adds an extra layer of security. If you lose a key or it becomes available to anyone else, without a password they can log in and impersonate you.

Also, most importantly, keep a copy of the private key in a secure (encrypted) backup. Treat that file as being as important as your passport.

Of course, all of this security hardening goes hand in hand with the requirement to keep the server's OS and *OpenSSH* software up to date by applying security patches and updates as soon as they become available.

In short, in any remote system, especially internet-connected ones, allowing username and password combos on the naked internet is just asking for trouble.

These are the bits that need to be changed. Ensure a backup copy first.

```
#AuthorizedPrincipalsFile none
#AuthorizedKeysCommand none
#AuthorizedKeysCommandUser nobody

# For this to work you will also need host keys in /etc/ssh
#HostbasedAuthentication no
# Change to yes if you don't trust ~/.ssh/known_hosts for
# HostbasedAuthentication
#IgnoreUserKnownHosts no
# Don't read the user's ~/.rhosts and ~/.shosts files
#IgnoreRhosts yes

# To disable tunneled clear text passwords, change to no here
#PasswordAuthentication yes
#PermitEmptyPasswords no

# Change to yes to enable challenge-response passwords (beware of
# some PAM modules and libraries)
```

Learning the way of the Docker

Docker is designed for an easy sysadmin life. Here are some top tips for using it.

Deploying new services can take a lot of time. Days sometimes. But *Docker* can make life a lot easier for system administrators. Containerisation can make deployment an extremely quick affair and management is simple. If the container goes wrong, delete it, redeploy, carry on.

To give an example, if an ELK server needs to be deployed, it's a bit of a chore, even with scripts. Deploying it with *Docker* takes less than 10 minutes from start to end. Doing it the 'old-fashioned' way takes a long time.

What makes using *Docker* even better is that all these images are self-contained.

Finished with it? Just **docker rm** it. Broken it somehow? Again, simply **docker rm** and then redeploy.

If it works on one *Docker* install, it will work on another. Laptop to server, cloud, K8s, no problem.

It takes some learning, but it is worthwhile, so here are a couple of useful tips for a happy *Docker* life:

- 1) Always use the **--name** feature.
- 2) Be careful how you attach to a host (**docker exec -it <container name> /bin/bash** for the win). Cancelling out incorrectly can shut down the container! **LXF**

CONTAINER ID	IMAGE	COMMAND	CREATED	STATUS
0b7f6b92ca67	bitnami/dokuwiki:latest	"/opt/bitnami/script..."	19 hours ago	Up 19 hours
0df0f14a81d7	phpmyadmin	"/docker-entrypoint..."	4 days ago	Up 2 days
4509623025e1	invidious-invidious	"/sbin/tini -- /invi..."	5 days ago	Up 2 days (healthy)
49a2b3a0303b	postgres:13	"docker-entrypoint.s..."	5 days ago	Up 2 days (healthy)

```
stu@workstation:~$
```

Using names with *Docker* makes things much easier to differentiate.



Stuart Burns is a Linux administrator for a Fortune 500 company specialising in Linux.

» LINUX NOT YET IN ARM'S REACH

This month, I have been doing a lot of *Docker* work, as you can probably tell from the article bottom-left. I kept thinking a Pi 5 with a true non-USB attached SSD drive would have made an excellent lightweight *Docker* host. The Pi Foundation teased that an NVMe add-on adaptor board is slated for early 2024 but it appears some people couldn't wait and have hacked a solution: <https://bit.ly/lxf310admin>.

Yes, it looks awkward and fiddly, but imagine a Pi 5 with 8GB and a 4TB NVMe disk. It would be a great portable mini server. Getting your hands on both items, though, would require luck and persistence.

The addition of SSD capability also opens up new uses for the Pi. Small media servers, bigger *Docker* images and not least better endurance over SD cards and better I/O throughput.

Before, space and write endurance were very much limiting factors. Myself? I am waiting for the official release because while the hack works, it's not production-ready or easy.

At the same time, Linux on ARM beyond SBCs seems very limited and rough and ready. Whilst it is straightforward to purchase an ARM-based server (AWS does them, larger hosting companies offer them as well), the laptop and desktop world have very little to talk of. Yes, there are some barely-more-powerful-than-a-budget-Android-phone types, but they just aren't up to the task.

With some ARM laptops starting to appear (for example, the ThinkPad X13s), the ARM ecosystem beyond SBCs and servers is yet to really kick off. Only time will tell...

Oxylabs

Mayank Sharma finds this proxy service to be a breath of fresh air.

IN BRIEF

With its high starting fees, bandwidth costs and not-for-newcomers nature, Oxylabs probably isn't for smaller projects. However, for everything else, it really is a premium service. Add to that the comprehensive list of multiple proxy types, and Oxylabs delivers for anyone who needs a premium quality service.

The Lithuania-based proxy service Oxylabs has been in the business since 2015. It has over 100 million proxies spread out in 195 countries around the world. On its website, the company takes pride in that they're ethically sourced "from vetted suppliers that are under strict obligations to inform end-users about their participation in a proxy network and acquire consent from them". Oxylabs focuses on meeting business needs (as opposed to stuff like sneaker proxies), such as brand protection, ad verification or SEO monitoring.

And while small providers leave you crossing your fingers and hoping they deliver what they promise, Oxylabs uses only certified data centres and upstream providers, and its plans are even insured by Lloyds (that's coverage for technology errors and omissions, and cyber insurance, if you're curious).

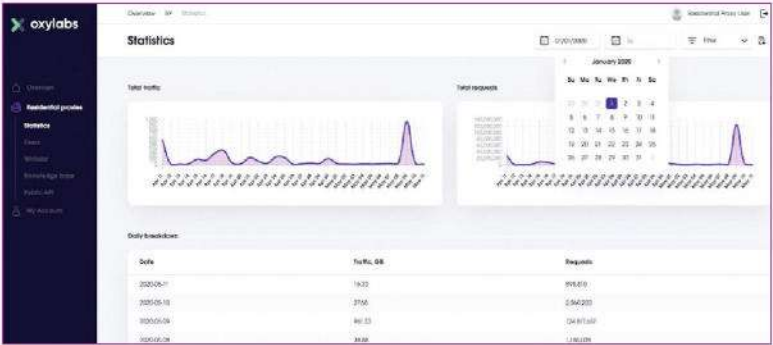
You can also opt for general-purpose search engine or ecommerce scraper APIs. There's even a Web Unblocker plan, which leverages AI and machine learning for a high success rate. Oxylabs claims it has "human-like browsing" to bypass CAPTCHAs, has adaptive HTML parsing, dynamic fingerprinting and JavaScript rendering to automatically retry if it runs into any problems.

Oxylabs' dedicated data centre proxies start at \$50 per month for your choice of 20 US data centre IPs, which works out at \$2.50 per IP. Residential proxy plans start at \$99 per month and offer 11GB of bandwidth, which comes to \$9 per GB. For those at the smaller end of the spectrum, there is also a Pay As You Go plan for the more occasional user. These are above average prices, but browse the feature list and it is understandable. Its products are built for large-scale projects, with few of the limits you'll find elsewhere. There are no restrictions on concurrent sessions, for instance, and there are no extra charges, either.

Geo-targeting enables you to work at the country, state or city level for maximum flexibility with your projects. Compare that to the extreme budget end of the market, where, for instance, some Storm Proxies plans have only two targeting options, US and EU, and you can see what you're paying for with Oxylabs.

Subnets have no set limits either. The company says that if you buy 1,000 proxies (for instance), you can generally have from four to 1,000 subnets, "depending on your business needs".

Furthermore, there's support for authenticating proxy access by username and password, or a trusted IP address. Some providers limit you to a set number of IPs, and on the lowest tiers, you're only allowed a single one. Oxylabs doesn't just allow you to add as



If you need a proxy service and no one else can help, call Oxylabs!

many specific IPs as you need, but you can even add ranges (192.168.1.0/24.)

However you set up your requests, Oxylabs claims excellent performance, with an average 99.95% success rate and 0.6 seconds proxy speed across its over 100 million residential IPs.

It's beyond the scope of this article to fully talk about Oxylabs' web scraper IPs, known as Web Unblocker, but they're just as comprehensive as the rest of the service. It can perform human-like browsing, and even get past CAPTCHAs. You can even decide where to store the data, if the standard S3 or Google Cloud storage isn't right; that's seriously flexible.

A knowledge base covers everything you need to know, from the dashboard basics to more advanced options to help customise the service for your precise needs. The Integrations page has useful tutorials on how to work with *Jarvee*, *Octoparse*, *Puppeteer*, *Proxifier* and many other third-party tools, for instance. Also, the more technical Developer's documentation is excellent, with high-quality articles and guides for using the service.

If you need more, though, Oxylabs offers live chat 24/7, and we got helpful responses within a couple of minutes. You can raise a ticket or send an email, but we couldn't find a direct phone number to just talk to your account manager. **LXF**

VERDICT

DEVELOPER: Oxylabs.io
WEB: https://oxylabs.io
PRICE: From \$10/GB PAYG (\$99 pm residential)

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	9/10	VALUE	7/10

Above average prices, but Oxylabs' huge network, wide range of products, smart web scraping tools and premium support more than justify the cost.

» Rating 8/10

Private Internet Access

A high-value VPN with a pile of interesting features that's keeping the nefarious **Mike Williams** out of trouble.

IN BRIEF

PIA is, and has been for a while, a solid contender in the VPN sphere – and it's easy to see why. You get tons of servers to choose from, it runs on just about any device, apps are easy to use and crammed with advanced features, and you'll have no trouble unblocking Netflix.

P rivate Internet Access (commonly known as PIA) is a capable VPN provider, now owned by Kape, which also owns CyberGhost, ZenMate and ExpressVPN. The company's network has speedy 10Gb/s servers spread across the world, including many areas that are often neglected by other providers. It recently expanded its network, passing from 84 to 91 countries. PIA told us that 30% of those servers are co-located (owned by the company rather than leased). This gives the VPN more control over the hardware and even improves speeds.

The service is also widely supported, with apps for Linux, Windows, Mac, Android and iOS, and browser extensions for *Chrome*, *Firefox* and *Opera*. Head to its website and you'll also find detailed setup tutorials for routers and other platforms, and smart DNS to unblock sites on consoles and other devices that can't run apps.

No limits!

Previously, PIA offered 10 simultaneous connections, but recently removed the limit entirely – now you can set up and use the service on as many devices as you need. That's far more generous than most of the competition.

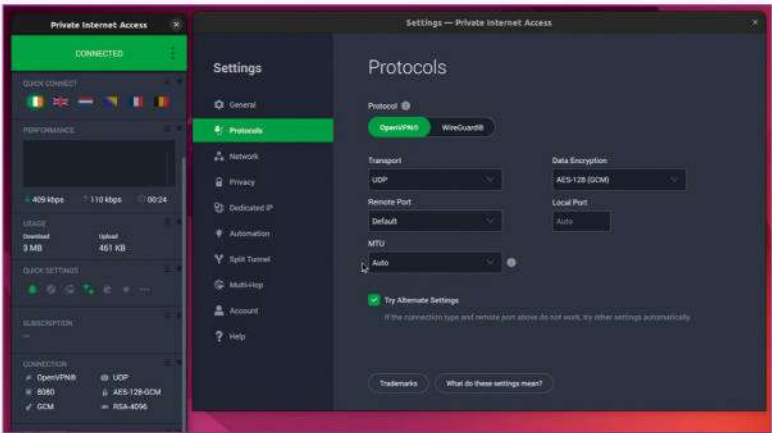
Extras range from built-in blocking of ads, trackers and known malicious websites to a SOCKS5 proxy for extra speed, port forwarding, preferred encryption, authentication, handshaking methods and more. If that sounds a little complicated, don't worry – there's 24/7 live chat support to talk you through the tricky bits.

The monthly plan is fair value at £9.99. The annual plan is cheaper than most at a low £2.71 a month, but the three-year plan drops to just £1.81. Pricing for a dedicated IP is reasonable at £4.50 a month, £3.83 on the one-year plan, and £2.25 over three years.

The apps are open source, enabling users to examine the source code, look for bugs, and see whether it's doing anything that might compromise their privacy. PIA also has a Bug Bounty programme. If a researcher finds and reports a genuine vulnerability, they could receive up to \$1,250, a bit on the stingy side but better than nothing.

PIA's Android app is certified by the ioXt Alliance, verifying that it complies with standards in areas such as cryptography, software verification and updates.

The company has a very clear logging policy, and states that the service does not store or share incoming and outgoing traffic information. This is backed up by an independent audit by Deloitte that interviewed staff, inspected its VPN and other servers, evaluated its policies, procedures and more. Deloitte's report found



It's good to see a fully-featured Linux interface on offer.

no evidence of privacy issues and backed up its no-logging claims.

Its *WireGuard* performance averaged 290-340Mb/s in the US and 350-360Mb/s in the UK. That's far behind the fastest VPNs – NordVPN, IPVanish, Surfshark and others hit over 950Mb/s. If you're setting up PIA on a router or other device, you may have to use OpenVPN, whose speeds reached 280-310Mb/s in the US. Surfshark's connections only managed 100-130Mb/s.

We tried to access the US (and other) libraries of Netflix, Amazon Prime, BBC iPlayer, Disney Plus, UK's ITV and Channel 4, and Australia's 9Now and 10Play from several test locations. It had no trouble getting into Netflix US, UK, Australia, Canada and Japan. PIA also defeated BBC iPlayer, ITV and Channel 4, giving it a 100% unblocking score. We saw no issues with torrents.

With a fully featured GUI control app for Ubuntu, Mint, Fedora, Arch and Debian, support is solid for Linux and it's simple to download and install. But if you get stuck, there's its Support Center with a web knowledgebase covering troubleshooting, account problems and more, alongside guides and live chat or email support. **LXF**

VERDICT


DEVELOPER: Private Internet Access
WEB: www.privateinternetaccess.com
PRICE: From £9.99 per month (£32.49 per year)

FEATURES	9/10	EASE OF USE	9/10
PERFORMANCE	8/10	VALUE	9/10

Speeds aren't the best, but otherwise this is a likeable VPN that gives you plenty for your money, plus it unblocked almost every streaming service we tried.

» **Rating 9/10**

CREDIT: Private Internet Access



No Ethernet, I expect you to die!

More stubborn than an anachronistic spy stereotype and just as reliable, *Darien Graham-Smith* uncrosses the network twisted pair.



We live in a wireless age. The latest Wi-Fi connections let us stream hundreds of megabytes per second through the air – shiny new Wi-Fi 7 routers are promising combined speeds of up to 40Gbits/s.

With all this data swirling around us, traditional cabled connections may seem old-fashioned to the point of obsolescence. But while Ethernet can't match the go-anywhere convenience of wireless networking, it still has several advantages over Wi-Fi.

While older readers might recall token-ring wired networks of yesteryear or the original 10BaseT form for Ethernet, network cables have not stood still. While wireless advancements have had all the headlines, our durable cable friends have steadily improved in speed.

With plans for 800GbE and 1.6TbE standards in the works, if you're contemplating laying cables for a wired network, we're here to outline what's coming, what's available and what'll be compatible in years to come.

Ethernet speeds

Exactly how fast is an Ethernet connection? The answer is more nuanced than you might expect. The Ethernet ports on most consumer devices use the IEEE 802.3ab standard, also known as 1000BASE-T, Gigabit Ethernet or just GbE. As the name implies, this supports a connection speed of 1Gbit/s, but that doesn't mean you can transfer data up and down the line at that rate. That's largely because, in addition to the data itself, Ethernet frames include routing information, plus checksum information to provide assurance that the data hasn't got garbled in transit.

Thankfully, this isn't a huge overhead. Gigabit Ethernet can convey around 920Mbits/s of data. That's much better than Wi-Fi, where real-world performance rarely approaches 50% of the speed on the box.

While Gigabit Ethernet is by far the most common implementation, it's not the only one. Older or cheaper gear may still have 100BASE-TX ports, offering only 10% of the bandwidth of a Gigabit connection (ironically, this is also known as Fast Ethernet). That might be plenty for a printer or a smart TV, but it can wreck the performance of a router or a NAS appliance.

This is where the simplicity of Ethernet can work against you. All the mainstream standards use the same cables and connectors, with transparent backward compatibility, so you can connect a Gigabit port to a Fast one without realising. The two devices happily proceed to communicate at 100Mbits/s, and you end up with a network connection that's underperforming for no obvious reason. It doesn't help that Ethernet ports are rarely labelled with their speeds; you might need to check an appliance's settings or look up the specs online to see if it's hobbled with a slower port.

Meanwhile, we're also starting to see high-end routers with faster-than-Gigabit Ethernet connectors. The 802.3bz-2016 specification defined new 2.5GBASE-T and 5GBASE-T standards, which enable communications at – you guessed it – 2.5Gbits/s and 5Gbits/s respectively. These technologies enable you to transfer or stream files within your home network at

» FIVE REASONS WHY ETHERNET IS BETTER THAN WIRELESS

1 You can't beat Ethernet for simplicity. Just run a cable from one socket to another and boom, the connection is made. Assuming the devices at either end are correctly configured, the link speed, addressing and routing should all be sorted out automatically.

You don't need to worry about keeping up with the latest high-end standards either; a £5 cable is all you need to get the fastest speeds. There's no confusion about different plugs and sockets, as with USB – with the exception of high-end enterprise-grade hardware, all Ethernet links use standard RJ45 (registered jack) connectors, with handy plastic latches to ensure the cables can't be accidentally yanked out.

2 Ethernet is more secure than Wi-Fi. Yes, the latest WPA3 wireless security protocol makes it all but impossible for an intruder to spy on your Wi-Fi connection, but not all devices support it – and even if a spy can't get into your wireless network, they can infer things from its name, signal strength, traffic patterns and so forth. With Ethernet, there's no way for anyone to snoop on your data, unless they can gain access to the physical cables and connectors.

3 Ethernet is almost always faster than wireless. The latest Wi-Fi routers provide high

speeds over a short distance, but carry your laptop into the next room and performance quickly drops off. A Gigabit Ethernet connection can reliably keep up a data rate of around 100MB/s over any length of cable, up to a huge 100 metres – and as we discuss in the main text, there are variants that can go far faster.

4 Ethernet provides a rock-solid, consistent connection, with only tiny variations in the data rate compared to wireless. Environmental interference is a non-issue, and you don't normally need to worry about network congestion either; while wireless routers have to divide up their bandwidth between any number of clients, every Ethernet connection represents a dedicated full-speed lane.

5 As well as higher headline speeds, Ethernet has much lower latency than wireless. On a home network, you might see a ping time of 5ms when pinging the router from a laptop connected over Wi-Fi; when you plug in an Ethernet cable, that plummets to 0.8ms. This doesn't make a huge difference for online gaming, as the latency of the connection to the remote server is a much bigger factor, but it can have a subtle effect on how snappy and responsive your system feels when accessing local networked resources.



Left: USB adapters are cheap and readily available. Right: The highest of speeds demands the latest cables; here's a sexy CAT8 cable.

» I'VE GOT THE POWER

Almost every consumer networking appliance comes with a little power brick. For businesses with a busy network infrastructure, all these power supplies add up to a lot of points of potential failure – and it's not always convenient to deliver mains power to every location.

The solution is the Power over Ethernet standard. Actually, this is a collection of standards; the original PoE spec was published nearly 20 years ago as IEEE 802.3af-2003, but today the most popular variants are 802.3at-2009 and 802.3bt-2018, aka PoE+ and PoE++.

The two systems work in the same way, allowing devices to be powered through a standard Ethernet socket, while sending and receiving data over the same cable. Any cable that conforms to Cat5 or greater can be used. PoE+ delivers up to 25.5W of power along with

Gigabit data speeds, while PoE++ can manage a maximum of 71.3W and is additionally compatible with 2.5GbE, 5GbE and 10GbE transmission modes.

What makes PoE so versatile is that the power element is optional – almost every device that supports PoE also has a separate socket for a standard barrel plug, so you could also connect it to a regular Ethernet router or switch. For large offices, the solution is usually a switch with powered ports and a management interface that lets you track and control power consumption.

For a small office or a home deployment, however, it may make more sense to invest in a PoE injector. This is a small, cheap, mains-powered box that simply forwards the network connection while adding power, so the downstream device can operate off a single cable.

lightning speeds, and to take full advantage of multi-Gigabit broadband as and when it becomes available.


Beyond this lies an extended range of Ethernet standards that go far beyond the needs of any household network. Many businesses working with big data sets and hosted applications use 10GbE, while data centres and supercomputing clusters have options ranging from 25GbE up to 400Gbits/s. Work is underway on finalising 800GbE and 1.6TbE standards.

These speeds push the limits of what a regular Ethernet cable can reliably convey, so fibre-optic cable is often used instead of copper wire. To accommodate this, enterprise gateways and switches come with large recessed sockets into which you can slot a connector module that handles the optical-to-digital conversion.

Which cables are able?

Apart from those high-end optical implementations, all Ethernet cables follow the same basic design, with a

Metres of the stuff! The base standard usually allows for 100m lengths of CAT5/6.



Latches ensure the cables can't be pulled out accidentally.

standard RJ45 connector at each end. Often the only visible difference between two cables is the casing colour and the plug; some cables also use a tangle-free flat design, rather than the traditional round format.

Internally, however, Ethernet cables are not all alike. They divide into eight categories (some of which have further variants), according to their recommended operating frequency and internal shielding. Since these aren't things you can determine visually, the category is usually printed on the outside of the cable.

The good news is there are only a few cable types you're likely to encounter. The first four are obsolete: Cat4 is only certified for connections up to 16Mbits/s, while Cat2 and Cat3 are even slower at 10Mbits/s and 4Mbits/s respectively. Cat1 isn't approved for data transmission at all – it's only for analogue phone calls.

This leaves Cat5 as the de-facto entry-level Ethernet cable, as it were. In fact, even this is officially deprecated, but you may still find Cat5 cables in use in older networks, or occasionally bundled with a cheap Fast Ethernet device. That's because these cables are designed for connections at up to 100Mbits/s.

For a guaranteed Gigabit Ethernet connection, you need the Cat5e variant. These cables support the same nominal 100MHz signal rate as Cat5, but they're manufactured to a higher standard to reduce noise and crosstalk inside the cable; this enables reliable transmission at Gigabit speeds, and even 2.5GbE.

If you try to run a Gigabit connection over a Cat5 cable, the results are unpredictable. If the cable is short and good quality, and there aren't many other cables nearby causing interference, it will probably work at the higher speed. If the cable isn't up to it, you may get inconsistent performance as data packets get corrupted and have to be resent – or the whole link might be unstable, resulting in disconnections.

The next step up from Cat5e is Cat6, which further tightens the requirements for shielding and signal quality. Most homes have no need for Cat6, but it's a popular choice for offices using high-end networks, or planning to do so in the future, as it supports 5GbE connections, as well as copper-based 10GbE.

Once you reach these speeds, you start to have to think about not only the construction of the cable, but also its length. The most common types of Ethernet cable are rated for up to 100m, but if you're using Cat6 for a 10Gbits/s connection, the maximum approved



Adding 10GbE to your PC? You'll want a PCIe adaptor.

length is 55m. To achieve the 100m reach, you need a Cat6A cable, which adds a new type of shielding to cut out the external noise that otherwise becomes a problem at very high transmission frequencies.

After that we're into the realms of the cables you'll probably never encounter. Cat7 mandates shielding for each individual pair of wires, as well as for the whole cable; this enables it to support 10GbE over a 100m cable, but in practice it's basically redundant because Cat6A already does that job. Cat8 supports 25GbE and 40GbE over copper cable, with a comparatively short maximum length of 36m.

Currently, there's no Cat9 standard, and it remains to be seen whether there ever will be. The fibre-optic connections that are already used by high-performance networks are faster, lighter and thinner than any copper cable, and there's no good reason why they shouldn't go mainstream – apart from all the existing devices that have embedded Ethernet ports and controllers.

Link aggregation

If you're wishing you could get more than Gigabit speeds out of your regular Ethernet network, there might be a way. The technology is called link aggregation – also known as network bonding or teaming – as defined by the IEEE's 802.3ad standard.

In practice, this simply means using two Ethernet ports at once, with two cables, to make a double connection between devices. It's commonly offered on routers and NAS appliances as a cheap, convenient way of doubling network bandwidth.

Link aggregation comes with a few catches, though. It needs compatible hardware at both ends; don't assume you can just use any two ports on a router or computer. For a desktop PC, you may need a dual-port PCIe Ethernet card that supports aggregation.

A bigger problem is that almost all routers we've seen only support 802.3ad on one pair of ports. This

» SWITCHING ROUTES

In most simple Ethernet setups, client devices plug directly into the router, or into a socket on a mesh station or extender. This is fine, until you run out of ports. Then you're stuck. However, it's easy to add more physical connectors to your network with a basic Ethernet switch.

A switch is a simple box with a row of Ethernet ports at the front. One of these needs to be connected to a port on the router, and then you can plug anything you like into the other ports. Through the combined

cleverness of IP and the Ethernet standards, the router can address and manage clients 'through' the switch as if they were directly connected. You can even daisy-chain switches together to distribute more ports over a wider area.

A typical office might use a 24-port managed switch – managed meaning that it has an integrated web interface for monitoring and managing traffic. For home use, you can get a basic unmanaged switch with five GbE sockets for under £15.

CABLE TIES

“Fibre-optic connections already used by high-performance networks are faster, lighter and thinner than any copper cable”

means you can (for example) have a 2GbE connection from your computer to the router, or a 2GbE connection from the router to the NAS – but not both. One way or another, therefore, the link between the PC and the storage is going to be bottlenecked at regular Gigabit speeds.

There's no really elegant way around this. You could buy an Ethernet switch that supports multiple pairs of aggregated ports, but these are uncommon and expensive. If your router supports Wi-Fi 6 or 6E, you may be able to get faster than Gigabit speeds by using a wireless link between the computer and router – but so far, we've only seen Gigabit-beating speeds from the most expensive hardware, and only over same-room connections.

Frankly, the best solution is to give up on link aggregation and see whether you can move up to 2.5GbE. A USB 2.5GbE adaptor can be had for around £30, which can be plugged into any laptop or desktop – and into many NAS appliances that lack native 2.5GbE hardware. Even if your router doesn't support multi-Gigabit Ethernet, you can get a 2.5GbE-compatible switch for around £80, enabling client devices to talk to each other at top speed, while maintaining a Gigabit connection to the router. **LG**



For 10GbE you're still looking at almost £60 per port.



Models such as the Zyxel MP105 bring 2.5GbE down to almost £15 per port.


tom's **HARDWARE**

GET ALL THE ESSENTIAL BREAKING NEWS FOR THE TECH ENTHUSIAST!

No matter if you're building a PC, buying a laptop or learning about robots, Tom's Hardware has all the comprehensive knowledge you need.




**Scan & Subscribe
for free!**



GET READY TO BECOME A TRUE PYTHON EXPERT

Make Python work for you with tutorials on coding with Django, Flask, Pygame and even more useful third-party frameworks.



Ordering is easy. Go online at:

magazinesdirect.com

Or get it from selected supermarkets & newsagents

HotPicks



Mayank Sharma

has spent so much time dreaming about a Pi 5 that he won't waste more waiting for his to arrive and will find new apps instead.

HiFile » Textadept » KooHa » Warehouse » Brave » Feh » Dosage » Bugdom 2 » OpenArena (Quake3e) » Linux Assistant » GridPlayer

FILE MANAGER

HiFile

Version: 0.9.9.5

Web: www.hifile.app

A file manager is an integral part of a desktop environment. And while most do an adequate job of handling files and folders, for a majority of users, they aren't everyone's cup of tea. *HiFile* is one of the growing list of third-party file managers that make up for the shortcomings of the default ones.

The app is available as an AppImage. Grab it from its website, and make it an executable via your existing file manager, or with `chmod +x`.

One of the highlights of *HiFile* is that it's lightning fast and lightweight. On first launch, you're asked to select some basic settings, such as language and colour theme, which you can change later.

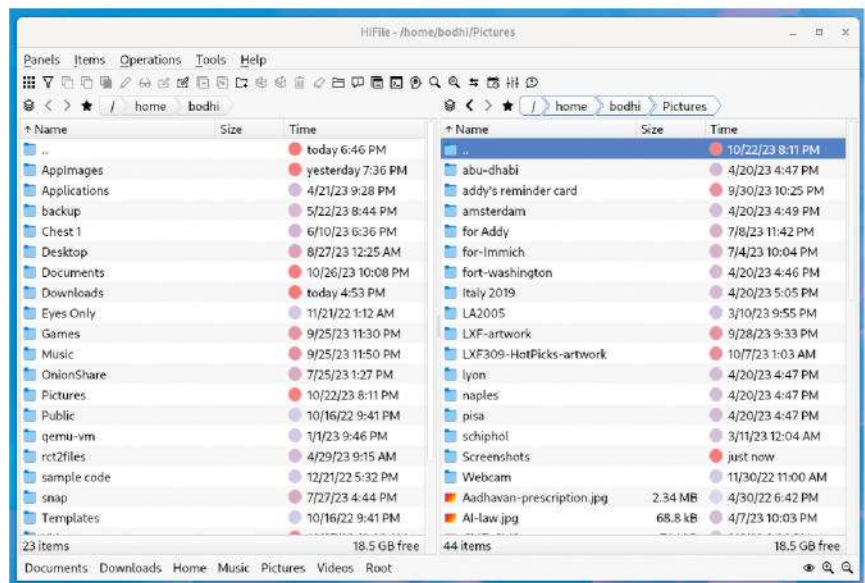
It uses a dual-pane layout, which is a traditional layout in many file managers. If you haven't used one before, you'll soon realise its convenience, especially when it comes copying and moving files and folders.

You can use *HiFile* for all common file manager tasks, such as renaming, deleting, creating archives or extracting items from them, viewing text files, binary files or images, and much more. Interestingly, *HiFile* treats archives as normal folders; you can open them and perform any operations as with normal folders.

You can also use it to find, filter and sort files or folders. If you have folders that you tend to visit often, you can bookmark them with `Ctrl+B` or the star icon. This then makes it easier to access them.

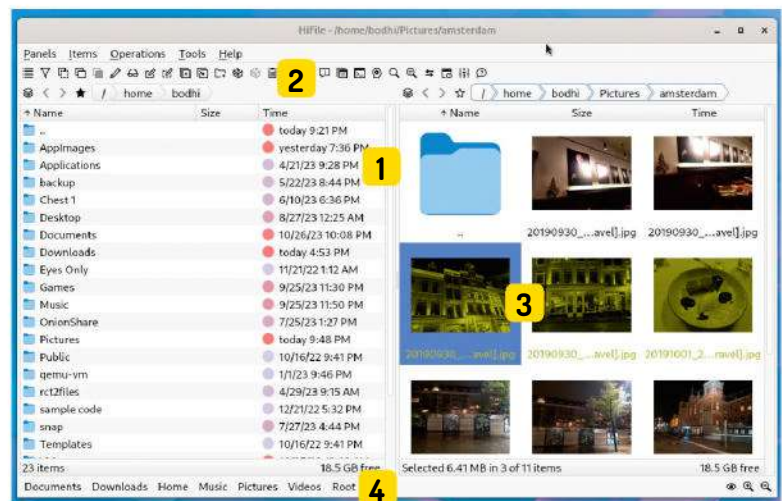
HiFile is highly customisable, too. You can also zoom in and out of the interface just as you do in a web browser. Hold down the `Ctrl` key and then either scroll with the mousepad or touchpad, or with the `+` and `-` keys.

Talking of keyboard controls, one of *HiFile*'s specialities is the use of the spacebar to jump from one location to another. Also, when you start typing the name of a folder, *HiFile* quickly points you to the correct folder.



HiFile has tons of useful features that make it a very efficient file manager, especially for keyboard warriors.

LET'S EXPLORE HIFILE...



1 Dual-pane

A two-pane view offers several advantages over a single-pane view, especially if you tend to find yourself juggling between multiple folders.

2 Toolbar

The toolbar lists all the commonly used functions, along with a tooltip to help new users get acquainted.

3 Item selection

HiFile shows selections in a different colour. Selections are also persistent, and aren't cancelled until you press `Escape` or leave the current folder.

4 Bookmarked folders

These are the bookmarked folders, which are also displayed at the top of all path searches for easier access.

TEXT EDITOR

Textadept

Version: 12.1 Web: <https://orbitalquark.github.io/textadept/>

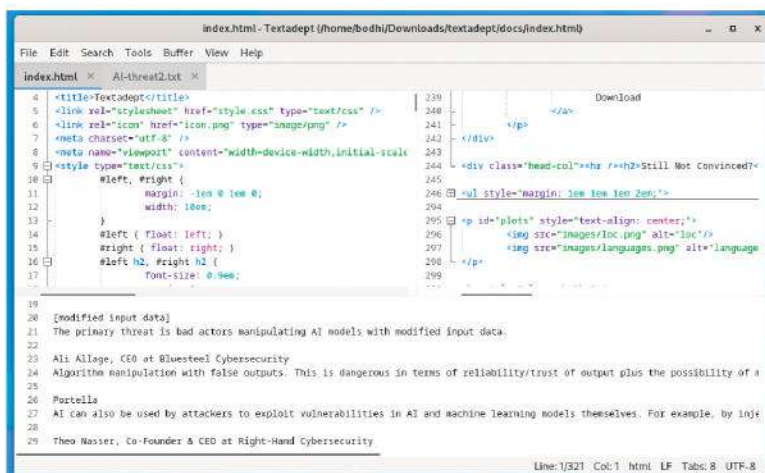
Most distros ship with a default text editor, which is usually quite vanilla. *Textadept's* philosophy is very similar to the venerable *Vim* in that it's an easy-to-use, yet highly extensible text editor, albeit with a GUI.

Textadept's binary package is self-contained and need not be installed. Just download it from the website and extract it with `tar zxvf textadept_12.1.linux.tgz`. Then change into the extracted directory, and fire up the editor with `./textadept`.

The app launches instantaneously, which is one of its main attractions. Its developers claim it takes up about 30MB of RAM and less than 10MB of disk space.

The interface is simple and optimised for speed and minimalism. While you can use it for regular text editing, it is particularly impressive as a code editor.

In addition to its tab support, the app supports unlimited vertical and horizontal view splitting, even of the same file. As well as its ability to autocomplete regular words, the editor can also autocomplete symbols for various programming languages. The app



supports over 100 different programming languages, and works with a plethora of formats and file types.

Despite being a graphical app, *Textadept* can be controlled entirely from the keyboard, which means nearly every menu item has a key binding. You can also reassign existing bindings or create new ones.

Interestingly, *Textadept's* menu is also accessible in the form of a searchable dialog that you can bring up with `Ctrl+P`. You can then type part of the name of any command in the dialog. Also impressive is the app's Find & Replace pane.

The app ships with a couple of themes to customise its appearance, and while it is pretty intuitive, there is a comprehensive manual on the website.

Textadept has a GTK version as well as a terminal version, which is ideal for editing on remote machines.

SCREENCASTER

Kooha

Version: 2.2.4 Web: <https://github.com/SeaDve/Kooha>


Gnome has a built-in screencaster and it works on Wayland. But it has no visible interface, nor any of the features you'd expect in standard screen-recording software. *Kooha* gives you the best of both worlds, being just as minimal as Gnome's built-in one, but with just enough interface to expose all the useful features.

The app is available on Flathub and can be installed with `flatpak install flathub io.github.seadve.Kooha`.

It has a basic widget-like interface, laced with easy-to-understand icons. Once you get used to it, you can access many of the functions via keyboard shortcuts.

Kooha can record either a section of the screen or the output of an entire screen or monitor. In either case, as soon as you hit the record button, the app asks you to point it to the monitor you wish to record from. You then get to select the area you want to record, if picked that option.

Before you hit the record button, you can ask the app to also capture the sound from the speakers, with or without the microphone input. By default, the app



The latest version of *Kooha* remembers the previously selected video source, such as a second monitor, which saves you a couple of clicks.

hides the pointer, but you can ask it to roll it in the recording as well.

In the app's minimal Preferences window, you can tweak the delay counter before you start recording, and change the location to save the recordings, which by default end up in `~/Videos/Kooha`.

Kooha encodes recordings in the WebM format by default, though you can choose to save them as MP4, MKV or even as a GIF. You can also change frames per second, if you know what you're doing. When it's recording, you get a simple counter, with a stop button.

FLATPAK MANAGER

Warehouse

Version: 1.3.0 Web: <https://github.com/flattool/warehouse>


If you've been trying our *HotPicks* recommendations of late, your Linux box must be chock-full of Flatpaks. *Warehouse* makes it fairly simple to manage Flatpak apps and has a number of useful Flatpak features that are typically only available from the CLI.

Not surprisingly, the app is available on Flathub, and can be installed with `flatpak install flathub io.github.flattool.Warehouse`.

When you fire up the app, it scans and displays an alphabetical list of all the installed Flatpaks on your box. You can use the search bar to filter through the list of installed Flatpaks, which comes in handy when you have as many Flatpak on your computer as we do.

When you find the Flatpak you're looking for, click on the corresponding info logo to view all kinds of useful details about the app, such as its version number, origin and installed size.

More functionality is wrapped under the three-dots menu adjacent to every app. There's the option to uninstall a Flatpak, browse the user data folder and



Use Warehouse's selection mode to either uninstall or delete the user data of a bunch of apps, or all of them, in a single click.

empty it, disable updates, or downgrade the app to a previous release.

One of *Warehouse*'s most interesting features is its ability to manage the application data of the installed Flatpaks. When you uninstall a Flatpak, some leave the user data behind, needlessly taking up space on your disk. You can reclaim this with *Warehouse*.

The app's hamburger menu includes a Manage Leftover Data option that scans the user data folder and ensures it is linked to an installed Flatpak. If it can't find an associated app, you have the option to either delete the data or attempt to install the matching app.

Manage Remotes is another useful option in the hamburger menu. You can use it to remove active Flatpak remotes and add new ones.

WEB BROWSER

Brave

Version: 1.59.124

Web: <https://brave.com>

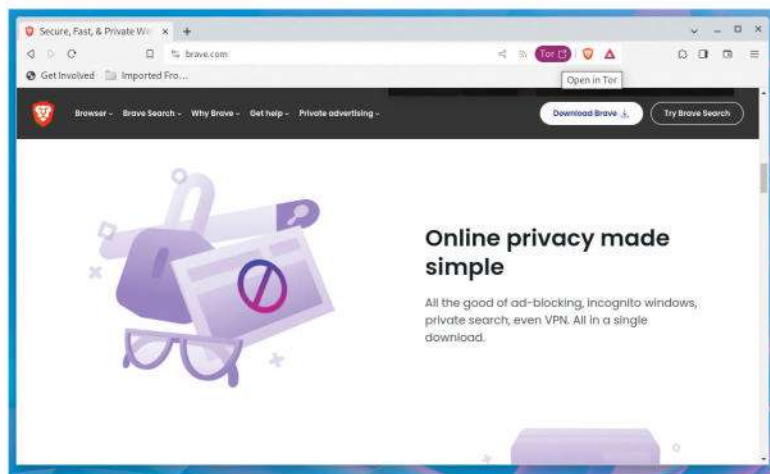
The web has gotten a lot nastier, and is a virtual minefield laced with privacy-invading cookies and trackers. While you can equip the popular web browsers to thwart unwanted information leaks, the *Brave* browser ships with privacy strengthening features built right into its core.

Brave has stable, beta and nightly repositories for both DEB and RPM-based distros, with step-by-step installation instructions available on its website.

On first run, *Brave* offers to import settings and bookmarks from other browsers. Its dashboard proudly displays the number of trackers and ads it has blocked, and the bandwidth and time this has saved.

Along with the list of recently visited websites, you can also ask the browser to enable the Brave News service. It's folly, in fact, to look at *Brave* as just another web browser. Instead, think of it as an ecosystem of apps and services.

One prominent component of that ecosystem is Brave Rewards, which again you can enable from the dashboard. When toggled, Brave Rewards displays the



occasional, unobtrusive advert. In return, if you view enough of them, you're rewarded in *Brave*'s native cryptocurrency, the Basic Attention Token. You can then pass these over to your favourite content creators.

Also, by default, the browser uses Brave Search, a privacy-centric search engine. You do, however, get the option to switch to any of the other options. Another interesting feature is its ability to display vertical tabs.

Besides the *Brave*-specific functionalities, the browser has all the features you'd expect from a mainstream web browser. As it's based on *Chromium*, the browser is compatible with all the extensions in the Chrome Web Store. Oh, and it hides those irritating cookie consent notices.

Brave ships with Tor, offers to open a 404 page using the Wayback Machine, and has built-in web page translation.

CLI IMAGE VIEWER

Feh


Version: 3.10.1

Web: <https://feh.finalrewind.org>

You don't have to be an experienced campaigner to know the usefulness of the command-line interface (CLI). Any task that you can do on Linux can be handled better on the CLI. Sure, while the GUI is convenient to operate, it can't match the speed and dexterity of the CLI. So, why should viewing images, a quintessential GUI task, be any different?

Feh is an image viewer that, unlike the ones you are used to, works on the CLI. And no, it doesn't display images inside the console as ASCII text. It shows them just like your everyday viewer does; it just lacks the GUI controls you get with other image viewers. Instead, it's completely keyboard-driven and can be set up with CLI arguments.

Feh doesn't put up binary packages, and is only available as a source tarball that you have to install manually. However, it's available in the official repos of most distributions, which saves you the effort of compiling it yourself. Users of Debian-based distros such as Ubuntu and Mint can install it with `sudo apt`



install *feh*, while `sudo dnf install feh` installs it for Fedora users.

To use *Feh*, just point it at the image you want to view. For example, `feh some-image.jpg` displays the image in a new window. The image is displayed at 100% resolution. You can use the up and down arrow keys to zoom in or zoom out, while the > and < keys can be used to rotate an image 90 degrees clockwise and anticlockwise. When you're done, press **q** to close the image and return to the CLI.

Instead of a single image, you can point *Feh* to an entire directory, such as `feh ~/Pictures`. This opens the first image, then you can use the right and left arrow keys to cycle through the entire folder.

As you'd expect, *Feh* has dozens of useful CLI options, and it's prudent to skim through its manual before using the viewer.

TREATMENT TRACKER

Dosage

Version: 1.1.1 Web: <https://github.com/diegopvlk/Dosage>

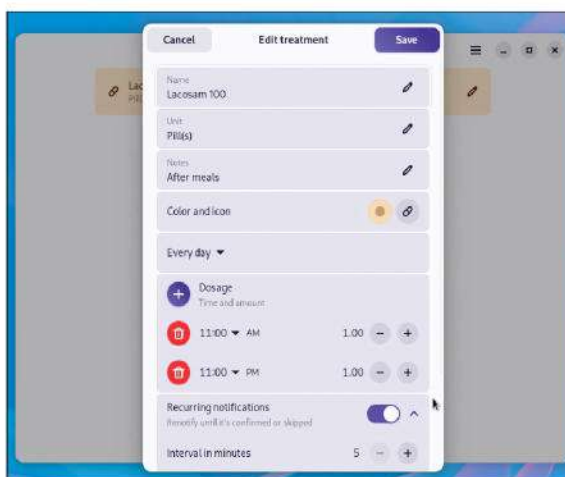
As we slip into the clutches of winter, we'll soon be lining our desks with bottles of pills. While it doesn't take much effort to repurpose any to-do app to remind us to take our medication on time, *Dosage* is purpose-built for this purpose.

Dosage is available on Flathub and can be installed with the following command: `flatpak install flathub io.github.diegopvlk.Dosage`.

The app has a clean interface. You begin by adding a treatment plan. This involves specifying the usual details, such as the name of the medication and whether it comes in the form of tablets, syrup, an injection or something else. You then specify the schedule of the medicine, whether you're supposed to take it every day, on specific days, or on a cycle.

Depending on the frequency, the app gives you further options to narrow down a schedule. For instance, if you select specific days, the app asks you to point it to the days you need to take the meds.

Once you have added the schedule, you need to enter details about the dosage, including the time and



For more personalisation, *Dosage* enables you to denote each treatment plan with a relevant colour and icon, such as a blue capsule or a yellow drop.

the amount required. This is particularly useful if you need to take different amounts of the medication at different times.

Besides these details, you can also optionally ask the app to renotify you about a dose until you tell it that you've taken the meds or skipped them on purpose. If you want, you can also optionally use the app to keep stock of your medication, and be mindful of the start and end of a treatment plan.

Dosage has enough dexterity to accommodate multiple medicines, with different schedules and potencies. The inventory and treatment management makes it very useful for setting up reminders for people with memory issues (or deadlines).

ACTION ADVENTURE

Bugdom 2

Version: 4.0.0 **Web:** <https://github.com/jorio/Bugdom2>

Bugdom 2 was originally released way back in 2002 by Pangea Software, and now it has been ported to Linux with the permission of the original developer.

The game is available on Flathub and can be installed with `flatpak install flathub io.jor.bugdom2` on top of any distro with Flatpak support.

The intro video sets up the game's premise. While on his way to visit his family, our protagonist, Skip the grasshopper, is ambushed by a Bully Bee that swoops down and steals Skip's knapsack. Our task as Skip is to track the Bully Bee and retrieve our knapsack.

Before you hit play, take a detour to the settings and look through the controls. By default, the game is set to Normal difficulty mode, but there's an easier Kiddie Mode on offer as well.

The chase takes place mostly around a garden. Skip can kick, jump and fly, and can also pick up stuff. These abilities come into play as you navigate the garden and interact with friend and foe. Among your friends are Sally the chipmunk, who gives you useful things, such



as maps, and marks checkpoints when you bring her acorns. Then there's Sam the snail who makes you do some tasks, such as finding his shell or freeing mice from traps, before helping you.

Enemies outnumber your friends and come in all shapes and sizes. To tackle them, you need power-ups, so keep your eyes peeled for them. These power-ups pop out when you run into butterflies, and you can use them to enhance your current abilities or get new ones.

In this initial release, the game has been given a 2023 upgrade. In addition to minor gameplay fixes, the developer has added support for widescreen monitors, polished the game's appearance, and improved support for input devices.

Although the game is over two decades old, Bugdom 2 is still engaging, more so if you haven't played the original.

FPS

OpenArena (Quake3e)

Version: 0.8.8

Web: www.openarena.ws

Another blast from the past. As long-time Linux gamers know, *OpenArena*, the once-popular *Quake III*-inspired first-person shooter, has been in a state of limbo for over a decade. While game development hasn't been renewed, a developer has put up a version of the game with the Quake3e engine.

Quake3e is pipped as a modern *Quake III Arena* engine that's faster, while still being compatible with all existing Q3A mods. More importantly, it has optimised OpenGL and Vulkan renderers and high-quality textures to give the graphics a much needed upgrade.

The game is available on Flathub and can be installed with the following command: `flatpak install flathub io.github.ec_Quake3e.OpenArena`.

Remember to head to the game's setup to configure the controls and customise your character. The gameplay is similar to the original and mirrors that



You can also play OpenArena online against other players, and the number of servers still available are testament to the game's popularity.

of *Quake III Arena*, with some differences. It can be played offline against bots and includes all the genre-defining gameplay modes, such as deathmatch, capture the flag, tournament, elimination, CTF elimination, last man standing, and more.

Click on a stage, choose your difficulty, and off you go. The objective is simple: frag as many of the opponents as possible using an assortment of weapons, from a shotgun to a rocket launcher, each with its own characteristics and use, depending on the situation.

OpenArea is fast-paced and takes some getting used to, even if you are familiar with the genre. And it's as violent as you'd expect any FPS to be. It doesn't ask much from your hardware, so you can run it without a hiccup on any computer that can run any of the modern Linux desktop environments. Sure, it doesn't look as neat as the newer FPS games, like *Xonotic*, but it's just as entertaining.

APP LAUNCHER

Linux Assistant

Version: 0.3.3

Web: www.linux-assistant.org

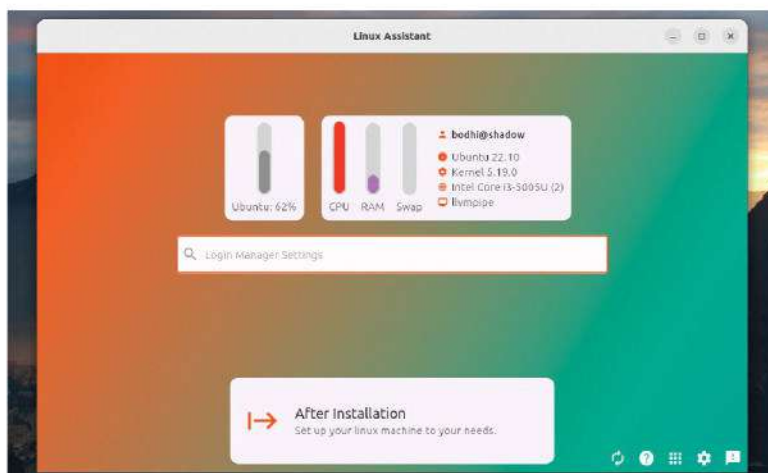
Linux desktops are far easier to navigate than they once were. That said, transitioning from Windows or Mac OS X can be jarring. *Linux Assistant* is designed to help new users find their way around Linux through a unified interface.

It works with Gnome, KDE, Cinnamon and Xfce desktops across most popular distros including Ubuntu, Pop!_OS, Mint and OpenSUSE. The only notable exception is Fedora, which isn't currently supported.

Grab the binary from its website and install it with your distro's package manager. On Ubuntu, for instance, install it with `sudo apt install linux-assistant.deb`.

When launched, the app takes you through a brief welcome process. It asks if it recognises your system correctly, then gets you to select a few default apps.

The app selection is part of the After Installation script, which you can invoke any time from within the app. Next, the script gives you two configuration



options. New Linux users should pick Automatic Configuration, which runs through common admin tasks, such as ensuring your installation is up to date, has all the proper multimedia codecs, and such.

The app essentially combines an app launcher with a very useful integrated search bar and a couple of handy scripts to ease administration tasks.

For instance, you can use *Linux Assistant* to launch apps or open folders you've used recently or bookmarked. Similarly, it can also help you get to recently used files, web pages and even browser bookmarks. In terms of administration, you can use the app to check your installation's security settings, set up automatic system snapshots, and more.

Linux Assistant's features vary on the mode of installation, as well as the distro and desktop environment you're running it on.

VIDEO PLAYER

GridPlayer

Version: 0.5.3 Web: <https://github.com/vzhd1701/gridplayer>

If the default video player in your distro doesn't work for you, there's always VLC. There's very little it can't do. That is, if you can navigate through its extensive menu, which is especially cumbersome for the average desktop user. This has given rise to a plethora of video players that wrap VLC's dexterity around an easy-to-use interface purpose-built for a specific task.

GridPlayer is one such VLC-based app that is designed to play multiple videos at the same time, the only limit being your hardware.

Download *GridPlayer's* Applmage from its website, and make it an executable with `chmod +x`.

The player launches with a bare interface that doesn't even have a menu bar. Instead, all the options are rolled up in the right-click context menu. Use it to open one or more video files, or just drag and drop them in the interface.

The player starts playing the videos when you open them. You can then keep on adding as many videos as you want, and the player keeps arranging them side





As its name suggests, *GridPlayer* is meant to play multiple videos in a grid. But you can ask it to zoom in and play a single video instead.

side in a grid. By default, all videos are muted, which is a good thing.

The video grid is configurable, and you can easily change its layout to suit your needs. For instance, you can rearrange the position of the videos by simply dragging and dropping them.

Use the spacebar to pause and resume all videos, though you can control playback for each video independently. Since it's based on VLC, *GridPlayer* can run pretty much every video format out there. Also just like VLC, it enables you to control various aspects of the playback, including the aspect ratio, playback speed, zoom and more.

In addition to offline videos, the app can also stream videos from YouTube and other websites because it includes `yt-dlp`. **LXF**



**Can you help inspire the
next generation of coders?**



Code Club is a nationwide network of volunteer-led after school clubs for children aged 9-11.

We're always looking for people with coding skills to volunteer to run a club at their local primary school, library or community centre for an hour a week.

You can team up with colleagues, a teacher will be there to support you and we provide all the materials you'll need to help get children excited about digital making.

There are loads of ways to get involved!
So to find out more, join us at **www.codeclub.org.uk**

CMAKE

Part One!
Don't miss
next issue,
subscribe on
page 16!

Write a Linux shell from scratch

Never one to shy away from the difficult, **Ferenc Deak** takes us by the hand and helps us code a shell – from scratch.



**OUR
EXPERT**

Ferenc Deak wanted to use Malbolge to create a Linux shell, but after several days in hell, he quickly came to his senses and continued the project in C++. Not that there is a huge difference...

In the Linux ecosystem, the shell is a vital component – a command-line interface (CLI) program, enabling users to interact with the underlying operating system through text-based commands. It acts as a bridge, facilitating communication between users and the core components of the operating system, which manage hardware resources and system services.

This series of tutorials embarks on a fascinating journey on how to write a Linux shell from the ground up, offering readers a unique opportunity to explore the inner workings of such a fundamental tool.

While shells can be written in scripting languages such as Python and Ruby, we're taking a more low-level approach for a modern spin. This article builds on the foundation of Chris Brown's article in **LXF149**, providing gentle guidance into programming a shell, with the twist of a fresh perspective. Switching from plain C to C++, we're using the **CMake** build system to streamline development and manage dependencies.

Due to space constraints, we're printing only the essential introductory code snippets. Subsequently, we'll reference different sections for a comprehensive understanding. To access the complete code for this shell, please visit <https://github.com/fritzone/lxf-shell>. We strongly encourage you to explore the code and build it yourself to gain a better perspective of how everything comes together (the code there has also comments that were removed here for the purpose of keeping the size of the tutorial manageable).

With all this in mind, we can embark on the perilous journey of writing a Linux shell. From scratch. The main reason is not that the world needs another shell, as there is already a plethora of them, but more to provide an opportunity to learn about the lower-level Linux system internals, to give you the chance to understand how stuff works behind the scenes, and while doing so, provide you with a basic introduction to build management with **CMake** (see boxout, page 93), and to

```
> $ ./lxf-shell
lxfsh$ ls
CMakeCache.txt CMakeFiles cmake_install.cmake lxf-shell Makefile
lxfsh$ ls -l -a
total 64
drwxr-xr-x 3 fld fld 4096 sep. 30 16:25 .
drwxr-xr-x 3 fld fld 4096 sep. 30 16:21 ..
-rw-rw-r-- 1 fld fld 14690 sep. 30 16:25 CMakeCache.txt
drwxr-xr-x 8 fld fld 4096 sep. 30 16:25 CMakeFiles
-rw-r--r-- 1 fld fld 1684 sep. 23 18:12 cmake_install.cmake
-rwxrwxr-x 1 fld fld 24440 sep. 30 16:25 lxf-shell
-rw-rw-r-- 1 fld fld 5287 sep. 30 16:25 Makefile
```

■ Showcasing the basic command execution via system, a wrapper for `/bin/sh`.

give a brief introduction to C++ as the choice of programming language of this guide.

Your first CMake

To start using **CMake**, you typically follow these steps:

1. Create a **CMakeLists.txt** file in your project's source directory to define project settings, source files, dependencies and executables (for now).
2. Create a **build** directory, to facilitate the out-of-source feature of **CMake** and to not to mess up your source tree, because **CMake** generates a lots of intermediary build files.
3. Run **CMake** to configure your project, by specifying the directory where your **CMakeLists.txt** file is, then optionally specifying a generator – under Linux, the default is plain Makefiles – and build configuration options if needed (with the **-D** flag) at a later stage.
4. Build your project using the generated build files, using your preferred build tool.

For now, we start with the following basic example:

```
cmake_minimum_required(VERSION 3.15)
project(lxf-shell VERSION 1.0)
set(CMAKE_CXX_STANDARD 20)
set(${PROJECT_NAME}_SOURCES
    ${CMAKE_CURRENT_SOURCE_DIR}/main.cpp
)
add_executable(${PROJECT_NAME} ${${PROJECT_NAME}_SOURCES})
```

The first **CMakeLists.txt** example is very straightforward: it sets minimum **CMake** requirements, creates a project (**lxf-shell**) and adds an executable,

QUICK TIP

Explore the code and build it yourself to gain a wider perspective on how everything comes together: <https://github.com/fritzone/lxf-shell>.

which will be again *lxf-shell*. The only tricky syntax is `${VARIABLE_NAME}`, which is used to retrieve the value of the given variable. For our case, the **project** keyword already set a variable called **PROJECT_NAME**, so we retrieve it and use it to create another variable, which after the substitution will be named **lxf-shell_SOURCES**. All good, so let's continue:

```
#include <iostream>
#include <cstdlib>
int main()
{
    while(true)
    {
        std::cout << "lxfsh$ ";
        std::string command;
        std::getline(std::cin, command);
        if(command == "exit")
        {
            exit(0);
        }
        system(command.c_str());
    }
}
```

Our first C++ example (call it **main.cpp**) is also very simple: it creates a command-line interface that continually prompts the user for input via **std::getline**. It executes the entered command using the **system** function, unless the user inputs **exit**, at which point the program terminates.

Save both files in the same directory. In the terminal, navigate to that directory. Once inside, create a **build** directory for the first iteration of the shell, and enter it:

```
$ mkdir build
$ cd build
```

Once in there, run **CMake** by specifying the directory where **CMakeLists.txt** is to generate the build files, then use **make** to build the project.

```
$ cmake ..
$ make
```

Running these two builds our current iteration of the shell, and from this point, executing your new shell is as simple as that seen in the screenshot (left).

And that's it, except there is just one issue. It does not behave like a proper shell, since the **system** call is just a clever wrapper around **fork** and **execl**, and uses the default shell, which should reside in **/bin/sh** of every decent Linux distro to execute the command.

Fork, knife and exec

In Linux programming, the **fork()** and **exec()** functions create new processes and start applications. These functions play a crucial role in process management and are essential for tasks such as running apps.

The **fork()** function is used to create a new process in Linux, but this does not mean that we've started a new application. When **fork()** is called, it duplicates the existing process, creating an identical copy known as the child process. The child process is a separate execution context with its own memory space, but it starts execution from the same point as the parent.

- **Parent process:** The process that calls **fork()**.
- **Child process:** The newly created process.

The **fork()** function returns different values to the parent and child processes. In the parent process, it returns the process ID (PID) of the child process, while

» WHAT THE SHELL?

Here are some key points we need to understand about the shell, and how these notions will be incorporated into our own shell.

1. Command Interpretation There are two categories of commands that we need to support:

- a. External commands (**ls**, **date**) are executable files that are already residing in the directories, and running them is the theme of the current tutorial.
- b. Support for built-in commands (**cd**, **alias**, etc) will be provided using the backbone plugin architecture we'll introduce soon.

2. Redirection and pipelines An introduction to developing the shell's support for features such as input/output redirection and pipelines will be started in our next instalment.

3. Shell prompt Customising the shell prompt will be an important part of our series. We will add features for displaying information such as hostname, current directory and more.

4. History and command-line editing Since we are unhappy with some features offered by current shell implementations concerning history, we will introduce a rather novel approach to how history is handled by our shell.

5. Tab completion Pressing the Tab key has a special place in our hearts and we will present an innovative approach to making this feature even more widely used by introducing a support library for command completion.

6. Customisation Implementing variables and aliases will be introduced very soon in this tutorial series, however support for functions will be delayed until we have the scripting feature in place.

7. Scripting Implementing a scripting language specific to our shell will be revisited at a later stage due to the sheer size of the project.

As a final touch, in homage to the magazine that accepted the publication of this author's blusterings, we shall call the shell *lxf-shell*.

in the child process, it returns 0. This allows the parent and child processes to distinguish themselves, and also the parent to have some control over the child.

The **exec()** function family is used to replace the current process image with a new one. After calling an **exec()** function, the existing program is terminated, and a new program is loaded into the same process space. This is commonly used to run external programs from within a C/C++ program.

At the current iteration of the shell we will use **execvp()**, which allows passing on a set of arguments to the application, and also searches for the program in directories listed in the **PATH** environment variable.

The **exec()** functions do not return if they are successful; they load the new program, and execution continues from there. If there's an error, they return -1.

Together, **fork()** and **exec()** are used for a range of tasks, from creating complex multiprocess applications to running different programs sequentially in a single process. They form the building blocks of process control and management in our favourite Linuxes. Armed with this knowledge, we can now create a more properly functioning shell, like the following:

```
#include <iostream>
#include <unistd.h>
```

```
> $ ./lxf-shell
lxfsh$ ls
CMakeCache.txt CMakeFiles cmake_install.cmake lxf-shell Makefile
lxfsh$ ls -l
execvp: No such file or directory
```

This happened when we just used **error:** the undisclosed failure of **execvp**.



```
> $ ./lxf-shell
lxfsh$ ls -l
execvp(pathname = "ls -l", argv = ["ls -l"]) failed, No such file or directory (
2, ENOENT) because there is no "ls -l" regular file in any of the command search
PATH directories ("/usr/local/sbin", "/usr/local/bin", "/usr/sbin", "/usr/bin",
"/sbin", "/bin", "/usr/games", "/usr/local/games", "/snap/bin")
```

When libexplain jumps to the rescue and provides meaningful description of the failure.

```
#include <sys/wait.h>
int execute(const std::string& program)
{
    pid_t child_pid;
    child_pid = fork();
    if (child_pid == -1)
    {
        std::perror("fork");
        return 1;
    }
    if (child_pid == 0)
    {
        const char* const arguments[] = {program.c_
str(), nullptr};
        if (execvp(program.c_str(), (char* const*)
arguments) == -1)
        {
            std::perror("execvp");
            return 1;
        }
    }
    else
    {
        int status;
        waitpid(child_pid, &status, 0);
    }
    return 0;
}

int main()
{
    while(true)
    {
        std::cout << "lxfsh$ ";
        std::string command;
        std::getline(std::cin, command);
        if (command == "exit")
        {
            exit(0);
        }
        execute(command);
    }
}
```

The **execute** function executes a given command by forking a new process. It starts by creating a child process using **fork()**. If the **fork** operation fails, it prints an error message and returns 1. In the child process (the true branch of the **child_pid == 0** check), it prepares the command and its arguments, and tries to execute it using **execvp()**. If **execvp()** encounters an issue, such as the command not being found, it prints an error message and returns 1 to indicate an error.

In the parent process, the function waits for the child process to complete using **waitpid()**. This ensures the parent doesn't continue executing until the child has finished. If the command executes successfully, the function returns 0. Now, if you compile and execute the program above, you are greeted with the familiar **lxfsh\$** prompt and you can again execute programs.

Oops, that didn't go as well as expected, did it? What is the problem with the **execvp** call when you

pass in a parameter to it? Fortunately, Linux comes with a handy library for such a complex situation as this. Its name is **libexplain** and it explains stuff to us, so the reasons for our failures will be written in a nice and understandable manner. If only life had a **libexplain**...

Explaining the unexplainable

To install the explainer library, run the following (or the equivalent for your distro):

```
$ sudo apt install libexplain-dev
```

Now we need to tell **CMake** to find an external library. Add the following lines in your **CMakeLists.txt**, just after the **add_executable** command:

```
find_library(LIB_EXPLAIN explain)
if(${LIB_EXPLAIN} STREQUAL "LIB_EXPLAIN-
NOTFOUND")
    message(FATAL_ERROR "You need to install
LibExplain to build this application")
else()
    message(STATUS "LibExplain found at: ${LIB_
EXPLAIN}")
endif()
target_link_libraries(${PROJECT_NAME} ${LIB_
EXPLAIN})
```

With these lines we have introduced several new commands for the aspiring **CMake** user:

- **find_library** tries to find a library. The first parameter for this invocation is the name of a variable that **CMake** stores the result in, while the second is the name of the library. In case **CMake** fails to find the library in question, the value of this variable is "**<VARNAME>-NOTFOUND**", so for our case it is "**LIB_EXPLAIN-NOTFOUND**". As **libexplain** is a standard Linux library, finding it requires none of the extra options that **find_library** can accept. Note how we didn't use the **lib** prefix when looking for **libexplain**, just plain **explain**? That's how the correct syntax is.
- **if** behaves like a classic **if** – it has a true branch from the **if** until the matching **else**, and a false branch from the **else** to **endif**. However, not all **ifs** must have an **else**. Please note how the value of the **LIB_EXPLAIN** variable is compared to the literal value of "**LIB_EXPLAIN-NOTFOUND**". This is how you compare strings in **CMake**... Ahem, did we mention that some developers find the syntax of **CMake** a peculiar beauty? Finally, the **endif** command closes the **if**.
- **message** prints a message to the screen. Depending on the optional first parameter – which in our case is **FATAL_ERROR**, used for really fatal situations, such as the inability to find a library, or **STATUS**, as in the situation of we just feel like saying something – the upcoming parameters are printed to the screen, and if there is not **FATAL_ERROR** for the first parameter, the execution just continues. Otherwise **CMake** stops processing and waits for the developer's remediation.
- **target_link_libraries** instructs **CMake** to generate code sequences that tell the linker to link our application with the libraries following.

With all these prerequisites in place, we can head to our **main.cpp**, and it's time to ask the newly introduced **libexplain** to explain what the problem with **execvp** is. Let's add the following to the beginning of **main.cpp**:

```
#include <libexplain/execvp.h>
```

In the body of the **execute** method, let's change the simplistic **std::perror("execvp")** to a more

sophisticated-looking:

```
std::cerr << explain_execvp(program.c_str(), (char*
const*)arguments) << std::endl;
```

Once you compile and run again, you can see in the screenshot (left) the response of the shell if you try again to convince it to run `ls -l`:

This is a lot more detailed, and it helps us identify what the real problem is: **execvp** thought we wanted to execute a program named `ls -l` instead of the program `ls` with the parameter `-l`, so obviously we have a few logic errors in our program.

Splitting up is easy to do

Very specifically, for our situation, trying to execute `ls -l`, **execvp** expects the following parameters:

```
char *program = "ls";           // Name of the program
char *arguments[] = {"ls", "-l", nullptr}; // Command-line arguments
```

However, we are providing instead:

```
char *program = "ls -l";        // Name of the program
char *arguments[] = {"ls -l", nullptr}; // Command-line arguments
```

And now we see why it fails. There is just one resolution to our problem: we need to split the incoming command string, and build up a decent **char*** array from it that **execvp** can work with. Splitting a string in C++ is not such a tedious task and it can be done in a few easy steps. A naive and straightforward implementation can be seen below:

```
std::vector<std::string> splitStringByWhitespace(const
std::string& input)
{
    std::vector<std::string> result;
    std::istringstream iss(input);
    std::string token;
    while (iss >> token)
    {
        result.push_back(token);
    }
    return result;
}
```

For the current iteration of the shell, this will do, and in order to integrate the newly acquired function, we shall modify our **execute** function to use it. The following piece of code goes into the `if (child_pid == 0)` branch – in other words, the place where we start our child process after a successful fork:

```
auto split = splitStringByWhitespace(program);
const char** arguments = new const char*[split.size() + 1];
size_t i = 0;
for( i < split.size(); i++) arguments[i] = split[i].c_str();
arguments[i] = nullptr;
if (execvp(arguments[0], (char* const*)(arguments)) ==
-1)
{
    std::cerr << explain_execvp(program.c_str(),
(char* const*)(arguments)) << std::endl;
    return 1;
}
```

This iteration uses a vector called **split** to store the command's components, converts them to C-style strings, and places them in an array of **const char***

» THE MAKING OF CMAKE

CMake is an open source, cross-platform build system. Over the years, it became the de facto build system for C++ developers who need to manage the build process for their software projects in a platform-independent and efficient manner.

Regardless of the fact that its syntax is considered ugly by some, **CMake** has gained widespread adoption in the C++ software development community due to its ability to simplify the build process, promote cross-platform development, and provide a unified approach to managing project configurations and dependencies.

CMake comes by default in the repositories of most Linux distros, so normally it should be enough to install it using the package management tool of your system. For Linux derivatives using **Apt** as their package management tool, use the following command (it should be very similar for **DNF** or **Zypper**, too):

```
$ sudo apt install cmake
```

CMake also is available in the Snap Store, so feel free to use that if you want the latest version:

```
$ sudo snap install cmake --classic
```

If you wish to overly complicate your life by balancing on the bleeding edge of the infamous double-edged sword of open source software development, you can download the latest sources from <https://cmake.org/download>, and compile for yourself. But since you took this path, we assume you have the necessary knowledge and will happily let you struggle with the process.

called **arguments**. A **nullptr** is added as the last element to terminate the argument list, and **execvp()** is called to execute the command. If **execvp()** fails, it prints an error message using a custom function (**explain_execvp()**) and returns an error code (1).

And with this in place, we can finally build our shell again, and after starting it, we can verify whether it accepts parameters to the commands as expected (see screenshot, below). And yes, it works – finally our commands can be properly parametrised.

A glimpse into the future

As we bring this article to a close, we're excited to offer a glimpse into what the next instalments have in store. Next month, we'll explore the intricacies of command redirection. These fundamental concepts are essential for understanding how shells can manipulate data streams, allowing for even more powerful and versatile command-line experiences.

So, be sure to stay tuned for the upcoming chapters in our journey. In the meantime, keep those coding spirits high, and happy coding, as you continue to explore the world of Linux programming! **LXF**

Properly executing commands via **execvp** yields the same result as system, but with better control for the programmer.

```
> $ ./lxf-shell
lxfsh$ ls -a -l --color
total 96
drwxrwxr-x 3 fld fld 4096 okt. 1 16:33 .
drwxrwxr-x 3 fld fld 4096 okt. 1 16:33 ..
-rw-rw-r-- 1 fld fld 14309 okt. 1 16:33 CMakeCache.txt
drwxrwxr-x 5 fld fld 4096 okt. 1 16:33 CMakeFiles
-rw-rw-r-- 1 fld fld 1682 okt. 1 16:33 cmake_install.cmake
-rwxrwxr-x 1 fld fld 55408 okt. 1 16:33 lxf-shell
-rw-rw-r-- 1 fld fld 5293 okt. 1 16:33 Makefile
```

» **DON'T MISSED THE NEXT PART!** Subscribe now at <http://bit.ly/LinuxFormat>

HOME ASSISTANT

Credit: www.home-assistant.io

Part three!
Don't miss
previous issues,
order on
page 64!

Build a smart-home data application

Matt Holder investigates how to take data from an API and display it in a GUI for fun and possibly profit!



**OUR
EXPERT**

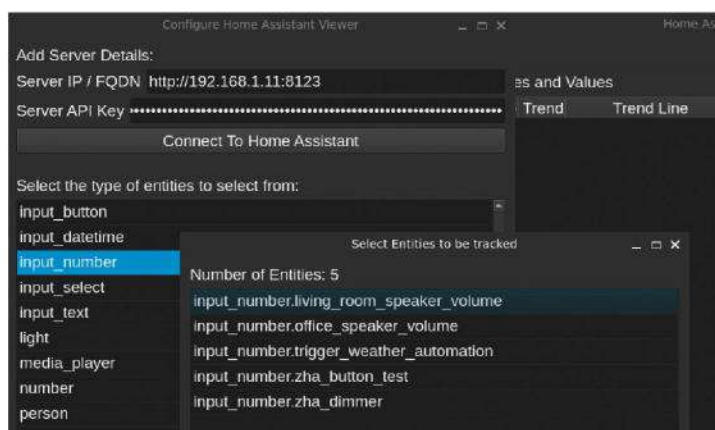
Matt Holder is an IT professional of 15 years, Linux user for over 20 years, user of plenty of home automation gear and self-professed geek.

We're putting the finishing touches to our smart-home data application. Our aim is to pull data from the *Home Assistant* API and then render it on screen in a GUI. We've split the program into two and in the first article we wrote the library that interacts with the API. Last month, we started writing the GUI application and we covered the entity selection window. This month, we will write the rest of the app, which consists of the configuration and main windows. As part of this article, we will be calling the API communication library, which will perform the actual data requests from the API.

Before we start coding again, let's have a look at *Home Assistant*. It is one of the most popular open source projects on GitHub and provides a way of linking together home automation equipment from different manufacturers, which potentially all use different protocols. *Home Assistant*'s aim is to keep as much data as possible within the internal network, but it can interact with many cloud services as well. Data gathered by *Home Assistant* is collected in what are called entities, and a history of entities is automatically kept and can be viewed from the web interface. *Home Assistant* provides an API that can be used to extract a list of entities within the system as well as the state of any entity. This app will enable you to select a number of entities to display and if possible a trend line will be displayed. The API is also capable of more functionality, but this is largely outside of the scope of this article.

Let's make a start by considering the configuration window. The code samples shown are only partial so that we can cover as many concepts as possible. The screenshot (opposite page) shows the configuration window, which will enable you to extrapolate the code shown with what is on the window itself.

```
class ConfigWindow(QMainWindow):
    def __init__(self, entityWindow, uri, apiKey,
                 windowWidth = 600, windowHeight = 500):
```



By the end of this article, we will have discussed the concepts required to build a GUI application like the one shown here.

```
super().__init__()
self.entityWindow = entityWindow
self.uri = uri
self.apiKey = apiKey
# CODE: Set the window's title and dimensions
# CODE: Create the widgets, add them to layouts
and display them on the screen
# CODE: Connect the signals to the slots
```

In this code sample, we define a class, called **ConfigWindow**, which is used to represent the configuration window. This new class inherits from the Qt-provided **QMainWindow** class. When initialised, the class has information relating to the **entityWindow**, the API URI and key, and default window size. The **super()** function is then used to initialise the methods provided by the base class, which means they are now available, too. Next, three instance variables are defined, which are used to store a reference to the **entityWindow** instance we have as well as the API details provided from the configuration file. As shown by a comment, the next thing we do is to set the window's title and dimensions.


The next two commented lines represent the code that is required to draw the widgets on the screen and to connect them to functions. As can be seen from the

QUICK TIP

The complete code listing can be downloaded from <https://github.com/mattmole/LXF308-309-API-GUI>.

screenshot (right), the configuration window is first of all made up of a number of labels, text fields to enter the API details, and a button, which is used to initiate the API connection. Once the connection has been successful, the table at the bottom of the window is populated using a list of all entities, as returned from the API. Once the widgets are shown on the screen, signals from the connection button and the entity type table are connected to functions. When the Connect button is selected, a connection needs to be made to the API and the **connectToApi** function carries this out. The second signal is from the entity type table – whenever a selection is made, the entity table on the entity window needs to be populated with all entities from the selected types.

```
def connectToApi(self):
    uriEntered = self.haServerAddressText.text()
    apiEntered = self.haApiKeyText.text()
    allEntities = HaEntityStatus(uriEntered,
    apiEntered)
    try:
        allEntities.readAllEntities()
        if allEntities.responseCode >= 200 and allEntities.
responseCode <= 400:
            global entitiesJson
            entitiesJson = allEntities.entities
            entitiesJson = sorted(entitiesJson, key=lambda
x:x['entity_id'])
            entityDomains = []
            for entity in entitiesJson:
                if "." in entity["entity_id"]:
                    entityDomain = entity["entity_id"].split('.')[0]
                    if entityDomain not in entityDomains:
                        entityDomains.append(entityDomain)
            # Set the size of the table and configure it
            self.entityTypeTable.
setRowCount(len(entityDomains))
            self.entityTypeTable.setColumnCount(1)
            # Populate the entity domain table
            for i in range(len(entityDomains)):
                self.entityTypeTable.setItem(i, 0, QTableWidgetItem(
entityDomains[i]))
            else:
                print ("Could not connect to the API")
                errorBox = CustomQMessageBox("API
Connection Error", "Could not connect to the API.
Please check the credentials.")
```



The configuration window consists of a number of widgets including labels, text boxes a button and a table.

```
except (requests.exceptions.InvalidURL, requests.
exceptions.ConnectionError):
```

```
errorBox = CustomQMessageBox("Connection
Error", "Invalid URL. Please check the details.")
```

The next definition in the code sample is the **connectToApi** function, which is connected to a signal from the API connection button. The first thing we do within the function is to take the information from the server and API key text boxes and assign them to variables. It is important to get the information from the text fields as it is not necessarily the case that the data in the configuration file is correct. Following this, we create another variable, called **allEntities** and assign to it an object from our API library, using the server address and API key from the previous variables. We now use a **try: except:** block to try running some code. If the code fails, the failure can be caught and processed accordingly. Within the **try** block, we first of all call the **readAllEntities()** method from the API library. As the name suggests, this method returns all entities registered in the connected-to *Home Assistant* installation. If the response code from the API is 200, or 201, which represents success, we can carry on using the information from the API.

We next define a global variable, called **entitiesJson**, which we can reference from across the code, and this stores a list of all entities that are returned from the API, which is in JSON format. To help with making the data as simple as possible to process, we use the **sorted** function to sort the entities into alphabetical

QUICK TIP

Why not add a tab to be able to send data back to *Home Assistant*? See this link for further details: <https://developers.home-assistant.io/docs/api/rest/>.

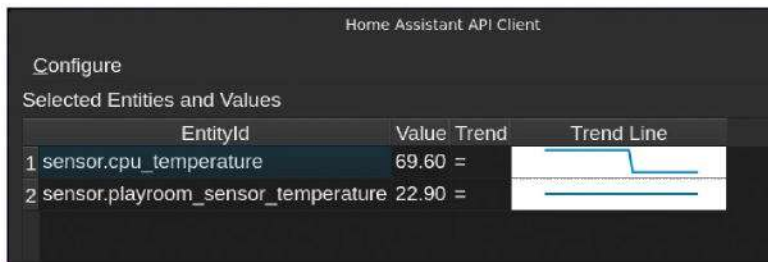
» GUI ZERO VS QT

GUI Zero is a wrapper on top of Python's *TKInter* graphical toolkit, which makes it simple to use, but it's not as powerful, extensible or nice looking as *Qt*. Replicate the entity window:

```
from guizero import App, Text, PushButton, Window,
MenuBar, TextBox, CheckBox, ListBox
app = App(title="HAClient")
entitiesWindow = Window(app, title="Select entities to
display", visible=False)
entitiesText = Text(entitiesWindow, width="fill", text="Number
of entities: 0")
entitiesListbox = ListBox(entitiesWindow, items=[],
height="fill", width="fill", multiselect=True, scrollbar=True,
command=updateEntityObjects)
```

```
entitiesListbox.enabled = False
```

In this code sample, we first of all import the necessary methods from the *guizero* library. On the second line, we create a variable and use the **App** method to create an app. When calling the function, we set the title. The next item we create in this code sample is the **entitiesWindow**, by calling the **Window** function. The first argument is the application that it relates to, followed by the title and whether the window is visible. In this instance, the window should not be visible because it will be made visible later in the process. On the next three lines, we place a text box and list box on the window, which can be used in a very similar way as the entities on the *Qt* application.



Here you can see the main window with two entities selected and displaying data.

QUICK TIP

You can learn about the Home Assistant API here: <https://developers.home-assistant.io/docs/api/rest/>.

order. What we do next is take each entity name and remove the entity domain or type from the front. Each entity begins with, for example, **sensor.** – the entity domain. To determine a list of domains, we loop through each entity stored in the **entitiesJson** variable and strip off the domain. If the calculated domain is not already stored in the **entityDomain** list, it is added. The entity domain table is then populated with the domains that have been calculated from the list of entities. The final part of this function deals with displaying message boxes if either the return code is not 200 or 201 or if the original call to the API method fails.

```
def selectEntityTypes(self):
    entityWindow.show()
    selectedDomains = set()
    for selected in configWindow.entityTypeTable.selectedRanges():
        for i in range(selected.topRow(), selected.bottomRow()+1):
            selectedDomains.add(configWindow.entityTypeTable.item(i,0).text())
    relevantEntitiesList = []
    counter = 0
    for entity in entitiesJson:
        for domain in selectedDomains:
            if entity["entity_id"].startswith(domain):
                relevantEntitiesList.append(entity["entity_id"])
            counter += 1
    self.entityWindow.entitiesTable.setRowCount(counter)
    self.entityWindow.entitiesTable.setColumnCount(1)
    self.entityWindow.numEntitiesLabel.setText("Number of Entities: " + str(len(relevantEntitiesList)))
    for i in range(0, len(relevantEntitiesList)):
```

```
self.entityWindow.entitiesTable.setItem(i, 0, QTableWidgetItem(relevantEntitiesList[i]))
```

The final function in the code sample is that of **selectEntityTypes**. The aim of this function is to detect which entity types have been selected in the entity types table and then populate the entities table in the entities window. The first thing we do in this function is store in a set called **selectedDomain** the names of entity types that have been selected (calendar, sensor, binary sensor and so on). Think of a set as being similar to a list, but duplicates cannot be present. Even if the same string is added twice, it does not return errors, but is only displayed once when the contents are accessed.

Now we loop through the list of selected ranges and determine which entity types have been selected. These are then added to the **selectedDomain** set. We now need to generate a list of suitable entities that can be displayed in the main window. A list, called **relevantEntitiesList**, is created, which is then used. To populate this list, we loop through all returned entities and within this, all **selectedDomains**. If the entity being selected starts with the domain selected within the inner loop, it is added to the **relevantEntitiesList**. Finally in this function, we loop through the **relevantEntitiesList** list and add these values to the table in the entity window.

Main window

Finally, let's discuss the **MainWindow** class and how we update the table, with updated values, every five seconds. What can be seen below is bare bones, with the trend line support removed as well as the code that relates to any concepts that have already been covered in previous sections of the series of articles. The screenshot (above-left) shows the main window, which will give you an idea of the widgets that have been used to make up the window. See the link in the Quick Tip (page 94) to download the complete code listing. Following this code sample, we will discuss everything that is printed.

```
class MainWindow(QMainWindow):
    def __init__(self, windowWidth = 800, windowHeight = 500, font=defaultFont):
        super().__init__()
        # Set window title and size
        # Instance variables to store useful information
        self.entityIdDict = {}
```

» WRITING BACK TO HOME ASSISTANT

As has been mentioned, the API is not read-only and instructions can be sent back to Home Assistant. Entities can have values set, services can be called, events can be fired and templates can be rendered. Templates are interesting as the Jinja templating language can be used to process data from Home Assistant and format it in many different ways.

```
from requests import post
url = "http://IP_ADDRESS:8123/api/services/light/turn_on"
headers = {"Authorization": "Bearer TOKEN"}
data = {"entity_id": "light.study_light"}
response = post(url, headers=headers, json=data)
print(response.text)
```

In the above code sample, we are sending a **POST** message to the Home Assistant API to turn on a light. Within Home Assistant, the light is called **study_light**. On the first line of the sample, we import the **post** method from the requests library. The URI of the API endpoint is then defined by the **url** variable. The **headers** variable is used to store the API token, which is used for authorisation. The **response** variable is then used to call the **post** method to send the request to the API and, as can be seen, the **authorisation** headers and data are sent within the request. Finally, the response is printed, which displays the list of states that have changed during the time the service was being called.


```
# Set a 5 second timer
self.checkThreadTimer = QTimer(self)
self.checkThreadTimer.setInterval(5000) #5
seconds
self.checkThreadTimer.timeout.connect(self.
updateTableValues)
self.checkThreadTimer.start()
# Set the menu bar up
menuBar = self.menuBar()
menuBar.setFont(font)
configureMenu = menuBar.
addMenu("&Configure")
configureMenu.setFont(font)
configApiMenuAction = configureMenu.
addAction("Configure API")
selectEntitiesMenuAction = configureMenu.
addAction("Select Entities")
#Connect signals to slots to show the other
windows when the menu options are clicked
configApiMenuAction.triggered.connect(self.
showConfigWindow)
selectEntitiesMenuAction.triggered.connect(self.
showSelectEntitiesWindow)
```

The first thing we do in this code sample is create a new class called **MainWindow**, which inherits from the **QMainWindow** class, which is provided by the **Qt** graphical toolkit. On the next line, we initialise the **MainWindow** class by using the **__init__** function. Arguments to the function are the width and height of the window as well as the font object that will be used to set the text of the menu. Next, we set the title and size of the window using built-in functions and create some instance variables to store information about our selected **entityIds**.

The next section of code is related to creating a timer, which we will use to update the main table every five seconds. First of all, we need to create a **QTimer** object, which we assign to the **self.checkThreadTimer** instance variable. Once the instance of the **QTimer** has been created, we use the **setInterval** function to ensure this timer is called every five seconds. So, what we have now is a timer that fires every five seconds, but doesn't do anything. We rectify this by using the **timeout.connect** function to link to the **self.updateTableValues** function. Once the **start** method has been called, the timer will fire every five seconds, which in turn calls the function to update the table.

The next six lines of code relate to creating the menu, which appears at the top of the window. First of all, we create a variable and use it to reference the window's **menubar** function. We then set the font of the menu before adding the configure menu, for which we again need to set the font. Items within the menu are referred to as actions and these are then added by using the **addAction** function. As with many other widgets in **Qt**, when the action is selected, this emits a signal, which needs to be connected to a slot (or

function). On the next two lines, we connect the two actions with the **showConfigWindow** and **showSelectEntitiesWindow** functions.

```
# CODE: Create the widgets and layouts and
display on the screen
def showConfigWindow(self):
    configWindow.show()
def showSelectEntitiesWindow(self):
    entityWindow.show()
def updateTableValues(self):
    domainPlotTypes = ["input_number", "input_text",
"number", "sensor"]
    if len(self.entityIdDict) > 0:
        self.entityTable.setRowCount(len(self.
entityIdDict))
        counter = 0
        # For each entity, pull the latest value and use it
        for entityId in self.entityIdDict:
            # CODE: This is the same code that we
            discussed last month
            # Update the table accordingly
            self.entityTable.setItem(counter, 0,
QTableWidgetItem(entityId))
            self.entityTable.setItem(counter, 1, QTableWidgetItem(
f'{entityObj['rowValue']}'))
            self.entityTable.setItem(counter, 2,
QTableWidgetItem(""))
            counter += 1
```

The next code sample, which is represented by a comment, is where we create the widgets to be displayed. This has been removed for space reasons but is incredibly similar to concepts discussed in the previous article of the series and when we discussed the config window, above.

Two functions are now defined, which complete the same action on different windows. These functions are used to open the configuration and entity selection windows. The chain of events for these functions to be called is first of all that the user selects the Configure > Configure API or Configure > Select Entities action. When the action is selected, a signal is emitted and this is connected to the relevant function.

We trust that you have enjoyed this coding series and it has helped to give you ideas of projects that you would like to complete on your own. **Qt** is an incredibly powerful environment and can be used from the most basic to the most complicated projects. If you wish to start a little simpler, then **GUI Zero** is a good choice. **LXF**

QUICK TIP

You can find out more about an alternative graphical toolkit at: <https://lawsie.github.io/guizero/>.

```
matt@pop-os ~> curl \
-H "Authorization: Bearer API_KEY \
-H "Content-Type: application/json" \
http://AP_ADDRESS:8123/api/states/sun.sun
{"entity_id":"sun.sun","state":"below_horizon","attributes":{"next_daw
67698+00:00","next_dusk":"2023-10-11T17:58:10.792528+00:00","next_midn
4+00:00","next_noon":"2023-10-11T11:55:41+00:00","next_rising":"2023-1
0","next_setting":"2023-10-11T17:23:19.268851+00:00","elevation":-26.6
g":false,"friendly_name":"Sun"},"last_changed":"2023-10-10T17:29:35.68
":"2023-10-10T20:19:27.478599+00:00","context":{"id":"01HCDKX4KPCRJ29Z
ull,"user_id":null}}
```

■ For troubleshooting, the venerable **curl** command can be used to query the API and return the data.

» **IMPROVE YOUR CODE SKILLS** Subscribe now at <http://bit.ly/LinuxFormat>

NEXT MONTH



BLOCK ALL ADVERTS!

Send your unwanted network traffic to the Pi-hole!
We set up homewide filtering and monitoring.

Build an analogue computer

Before there was digital computing, analogue had to run things – we go back in time to build our own analogue systems.

Get a better file manager

You can do better than the default Gnome file manager; we test the best file commanders the open source world can offer.

Ultimate CCTV system

Create your own fully remote-controlled CCTV system with the best *Motion* software to monitor and protect your home.

The rise and fall of Basic

It's not gone or forgotten, but we chart the rise, fall and return of Basic to see where it came from and what it influenced.

Content of future issues subject to change. The Pi-hole might have claimed another victim...

LINUX
The #1 open source mag **FORMAT**

Future Publishing Limited,
Quay House, The Ambury, Bath, BA1 1UA
Email contact@linuxformat.com

EDITORIAL

Editor-in-chief Neil Mohr
Art editor Fraser McDermott
Production editor Katharine Davies
Group editor-in-chief Graham Barlow
Group art director Warren Brown
Editorial contributors
Denise Bertacchi, David Bolton, Neil Bothwick,
Stuart Burns, Christian Cawley, Ferenc Deak,
Nate Drake, Darien Graham-Smith, Matt Holder,
Jon Masters, Nick Peers, Les Pounder, Michael Reed,
Jacob Ridley, Mayank Sharma, Shashank Sharma,
Rich Stanton, Mike Williams
Cover illustration Magictorch.com
Raspberry Pi is a trademark of the Raspberry Pi Foundation.
Tux credit: Larry Ewing (lewing@isc.tamu.edu).
Ubuntu is a trademark of Canonical Limited. We are not
endorsed by or affiliated with Canonical Limited or the Ubuntu
project. The Tails logo is released under the GNU GPL v3.
Content production Adequate Media Limited

ADVERTISING

Commercial sales director Clare Dove
clare.dove@futurenet.com
Advertising director Lara Jaggon
lara.jaggon@futurenet.com
Account director Andrew Tilbury
andrew.tilbury@futurenet.com

INTERNATIONAL LICENSING

Head of print licensing Rachel Shaw
Linux Format is available for licensing and syndication.
To find our more contact us at licensing@futurenet.com
or view our content at www.futurecontenthub.com

NEW SUBSCRIPTIONS & BACK ISSUES

Web www.magazinesdirect.com
UK 0330 333 1113 World +44 (0) 330 333 1113

EXISTING SUBSCRIPTIONS

Web www.mymagazine.co.uk
UK 0330 333 4333 World +44 (0) 330 333 4333
Subscription delays: Disruption remains within UK
and international delivery networks. Please allow up
to seven days before contacting us about a late
delivery to help@magazinesdirect.com

CIRCULATION

Newstrade & retail category director Ben Oakden

PRODUCTION AND DISTRIBUTION

Group head of production Mark Constance
Production manager Nola Cokely
Senior ad production manager Jo Crosby
Digital editions manager Jason Hudson

THE MANAGEMENT

Managing director technology group Paul Newman
Global head of design Rodney Dive
Commercial finance director Tania Brunning
Printed by Wyndeham Peterborough, Storey's Bar
Road, Peterborough, Cambridgeshire, PE1 5YS
Distributed by Marketforce, 5 Churchill Place,
Canary Wharf, London, E14 5HU
www.marketforce.co.uk For enquiries email:
mfcommunications@futurenet.com

Linux® is the registered trademark of Linus Torvalds in the US and other countries.
GNU/Linux is abbreviated to Linux throughout for brevity. Where applicable,
code printed in this magazine is licensed under the GNU GPL v2 or later. See
www.gnu.org/copyleft/gpl.html. All copyrights and trademarks are recognised
and respected.

Disclaimer All contents © 2023 Future Publishing Limited or published under
licence. All rights reserved. No part of this magazine may be used, stored,
transmitted or reproduced in any way without the prior written permission of the
publisher. Future Publishing Limited (company number 2008885) is registered in
England and Wales. Registered office: Quay House, The Ambury, Bath BA1 1UA. All
information contained in this publication is for information only and is, as far as we
are aware, correct at the time of going to press. Future Publishing Limited cannot
accept any responsibility for errors or inaccuracies in such information. You are
advised to contact manufacturers and retailers directly with regard to the price of
products/services referred to in this publication. Apps and websites mentioned in
this publication are not under our control. We are not responsible for their contents
or any other changes or updates to them. This magazine is fully independent and
not affiliated in any way with the companies mentioned herein.

If you submit material to us, you warrant that you own the material and/or have
the necessary rights/permissions to supply the material and you automatically
grant Future Publishing Limited and its licensees a licence to publish your
submission in whole or in part in any/all issues and/or editions of publications, in any
format published worldwide and on associated websites, social media channels
and associated products. Any material you submit is sent at your own risk and,
although every care is taken, neither Future nor its employees, agents,
subcontractors or licensees shall be liable for loss or damage. We assume all
unsolicited material is for publication unless otherwise stated, and reserve the right
to edit, amend and adapt all submissions. All contents in this magazine are used at
your own risk. We accept no liability for any loss of data or damage to your systems,
peripherals or software through the use of any guide.

We are committed to only using magazine paper derived from
responsibly managed, certified forestry and chlorine-free
manufacture. The paper in this magazine was sourced and
produced from sustainable managed forests, conforming to
strict environmental and socioeconomic standards.



Future is an award-winning international media group and
leading digital business. We reach more than 57 million
international consumers a month and create world-class
content and advertising solutions for passionate
consumers online, on tablet & smartphone and in print.

FUTURE

Connectors.
Creators.
Experience
Makers.

Future plc is a public
company quoted on the
London Stock Exchange
(symbol: FUTR)
www.futureplc.com


Chief Executive Officer **Jon Steinberg**
Non-Executive Chairman **Richard Huntingford**
Chief Financial and Strategy Officer **Penny Ladkin-Brand**

Tel +44 (0)225 442 2444

CREDIT: Magictorch



For press freedom
with responsibility
www.linuxformat.com



The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

EFF.ORG

ELECTRONIC FRONTIER FOUNDATION

Protecting Rights and Promoting Freedom on the Electronic Frontier

IF I HAD SOMEWHERE TO LIVE...

I COULD GO ANYWHERE IN LIFE



When Abi's mum died, life got tough. She didn't get on with her dad and the arguments became violent. Abi felt her only choice was to leave home. With just the clothes on her back, and no idea where to go, she ended up sleeping on the streets in the freezing cold.

Right now, you could give a homeless young person like Abi somewhere to start their future

Abi's life changed when she was given a room at Centrepoint. A safe place to sleep and recover. A place to develop the skills and confidence she needed to rebuild her life – and leave homelessness behind for good. Now, Abi believes she can go anywhere.

Thousands of homeless young people like Abi are desperately trying to find their place in the world – but first they need a place to start again. **You could help right now by sponsoring a room at Centrepoint for just 40p a day.**

We know this support changes lives. 88% of the young people we help move on positively in life. So please, help someone like Abi today. **Thank you.**

Text PLACE to 78866 to donate £3

Call free on 0800 472 5798

Visit centrepoint.org.uk/place

Or complete and return the form below

SPONSOR A ROOM. HELP A HOMELESS YOUNG PERSON FIND THEIR PLACE.

YES, I WANT TO SPONSOR A ROOM FOR £12 A MONTH

Please collect my payment on the 1st/15th of every month (please circle preferred date).

Instruction to your Bank or Building Society to pay by Direct Debit 

To the Manager:
Name and full address of your
Bank or Building Society: _____
Postcode: _____

Originators Identification No. **659107**

Name(s) of Account Holder(s) _____

Bank Sort Code: -- Account Number:

Instructions to your Bank or Building Society: Please pay Centrepoint Direct Debits from the account detailed in this instruction, subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Centrepoint and, if so, details will be passed electronically to my Bank/Building Society.

Signature(s) _____ Date _____

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

Increase your donation by 25p for every £1 you donate with Gift Aid *giftaid it*

☐ By ticking this box I confirm I am a UK taxpayer and want Centrepoint to Gift Aid all donations I've made in the last four years and any donations I make in the future until I notify you otherwise. I understand that if I pay less Income Tax and/or Capital Gains Tax in any tax year than the amount of Gift Aid claimed on all my donations it is my responsibility to pay any difference.

Your name and address are needed to identify you as a current UK taxpayer.

Full Name: _____

Address: _____

Postcode: _____ Telephone: _____

Email: _____

Thank you for donating to Centrepoint today.

We'd love to show the impact of your support and share how young people continue to need your help, through newsletters, fundraising appeals and information about events. Please let us know how best to stay in touch with you by adding your details above, and ticking the relevant boxes:

☐ **Please contact me by email** ☐ **Please contact me by phone**

Last year our supporters helped us change the lives of over 15,000 homeless young people. By letting us know we can count on you, we can continue helping young people with a home and a future.

Your privacy is key.

We promise never to sell or swap your details, and will always keep them secure. You can view our Privacy Notice in full at www.centrepoint.org.uk/privacy. You can opt out of post and change how we communicate with you at any time. Please call **0800 232320** and speak to one of our friendly team or email supportercare@centrepoint.org to do this and ask us any questions.

Please return this form to: Freepost Plus RTXY-LBEA-UTJJ, Angel Fulfilment Services Ltd, Communication Centre, Par Moor Road, Par PL24 2SQ.

Registered office address: Central House, 25 Camperdown Street, London E1 8DZ.
Company no. 1929421. Registered Charity no. 292411. Housing association no. H1869.
VAT registration no. 649 345 018.

